

Unambiguous Prioritized Repairing of Databases*

Benny Kimelfeld[†]
Technion
Haifa 32000, Israel
bennyk@cs.technion.ac.il

Ester Livshits
Technion
Haifa 32000, Israel
esterliv@cs.technion.ac.il

Liat Peterfreund
Technion
Haifa 32000, Israel
liatpf@cs.technion.ac.il

ABSTRACT

In its traditional definition, a repair of an inconsistent database is a consistent database that differs from the inconsistent one in a “minimal way.” Often, repairs are not equally legitimate, as it is desired to prefer one over another; for example, one fact is regarded more reliable than another, or a more recent fact should be preferred to an earlier one. Motivated by these considerations, researchers have introduced and investigated the framework of preferred repairs, in the context of denial constraints and subset repairs. There, a priority relation between facts is lifted towards a priority relation between consistent databases, and repairs are restricted to the ones that are optimal in the lifted sense. Three notions of lifting (and optimal repairs) have been proposed: Pareto, global, and completion.

In this paper we investigate the complexity of deciding whether the priority relation suffices to clean the database unambiguously, or in other words, whether there is exactly one optimal repair. We show that the different lifting semantics entail highly different complexities. Under Pareto optimality, the problem is coNP-complete, in data complexity, for every set of functional dependencies (FDs), except for the tractable case of (equivalence to) one FD per relation. Under global optimality, one FD per relation is still tractable, but we establish Π_2^P -completeness for a relation with two FDs. In contrast, under completion optimality the problem is solvable in polynomial time for every set of FDs. In fact, we present a polynomial-time algorithm for arbitrary conflict hypergraphs. We further show that under a general assumption of transitivity, this algorithm solves the problem even for global optimality. The algorithm is extremely simple, but its proof of correctness is quite intricate.

1. INTRODUCTION

Managing database inconsistency has received a lot of attention in the past two decades. Inconsistency arises for different reasons and in different applications. For example, in common applications of Big Data, information is obtained from imprecise sources (e.g., social encyclopedias or social networks) via imprecise procedures (e.g., natural-language processing). It may also arise when integrating conflicting

data from different sources (each of which can be consistent). Arenas, Bertossi and Chomicki [3] introduced a principled approach to managing of inconsistency, via the notions of *repairs* and *consistent query answering*. Informally, a *repair* of an inconsistent database I is a consistent database J that differs from I in a “minimal” way, where *minimality* refers to the *symmetric difference*. In the case of anti-symmetric integrity constraints (e.g., denial constraints and the special case of functional dependencies), such a repair is a *subset repair* (i.e., J is a consistent subinstance of I that is not properly contained in any consistent subinstance of I).

Various computational problems around database repairs have been extensively investigated. Most studied is the problem of computing the *consistent answers* of a query q on an inconsistent database I ; these are the tuples in the intersection $\bigcap \{q(J) : J \text{ is a repair of } I\}$ [3, 29]. Hence, in this approach inconsistency is handled at query time by returning the tuples that are guaranteed to be in the result no matter which repair is selected. Another well studied question is that of *repair checking* [1]: given instances I and J , determine whether J is a repair of I . Depending on the type of repairs and the type of integrity constraints, these problems may vary from tractable to highly intractable complexity classes. See [4] for an overview of results.

In the above framework, all repairs of a given database instance are taken into account, and they are treated on a par with each other. There are situations, however, in which it is natural to prefer one repair over another [8, 16, 36, 37]. For example, this is the case if one source is regarded to be more reliable than another (e.g., enterprise data vs. Internet harvesting, precise vs. imprecise sensing equipment, etc.) or if available timestamp information implies that a more recent fact should be preferred over an earlier fact. Recency may be implied not only by timestamps, but also by evolution semantics; for example, “divorced” is likely to be more updated than “single,” and similarly is “Sergeant” compared to “Private.”¹ Motivated by these considerations, Staworko, Chomicki and Marcinkowski [36, 37] introduced the framework of *preferred repairs*. The main characteristic of this framework is that it uses a *priority* relation between conflicting facts of an inconsistent database to define a notion of *preferred repairs*.

Specifically, the notion of *Pareto optimality* and that of *global optimality* are based on two different notions of *improvement*—the property of one consistent subinstance being preferred to another. Improvements are basically lifting

*This research is supported by the Israeli Science Foundation, Grant #1295/15.

[†]Taub Fellow, supported by the Taub Foundation.

¹See [15] for a survey on aspects of data quality.

of the priority relation from facts to consistent subinstances; J is an improvement of K if J contains a fact that is better than all those in $K \setminus J$ (in the Pareto semantics), or if for every fact in $K \setminus J$ there exists a better fact in $J \setminus K$ (in the global semantics). In each of the two semantics, an *optimal repair* is a repair that cannot be improved. A third semantics proposed by Staworko et al. [37] is that of a *completion-optimal* repair, which is a globally optimal repair under some extension of the priority relation into a *total* relation. In this paper, we refer to these preferred repairs as *p-repair*, *g-repair* and *c-repair*, respectively.

Fagin et al. [13] have built on the concept of preferred repairs (in conjunction with the framework of *document spanners* [14]) to devise a language for declaring *inconsistency cleaning* in text information-extraction systems. They have shown there that preferred repairs capture ad-hoc cleaning operations and strategies of some prominent existing systems for text analytics [2, 9].

Staworko et al. [37] have proved several results about preferred repairs. For example, every c-repair is also a g-repair, and every g-repair is also a p-repair. They also showed that p-repair and c-repair checking are solvable in polynomial time (under data complexity) when constraints are given as denial constraints, and that there is a set of functional dependencies (FDs) for which g-repair checking is coNP-complete. Later, Fagin et al. [12] extended that hardness result to a full dichotomy in complexity over all sets of FDs: g-repair checking is solvable in polynomial time whenever the set of FDs is equivalent to a single FD or two key constraints per relation; in every other case, the problem is coNP-complete.

While the classic complexity problems studied in the theory of repairs include repair checking and consistent query answering, the presence of repairs gives rise to the *cleaning problem*, which Staworko et al. [37] refer to as *categoricity*: determine whether the provided priority relation suffices to clean the database unambiguously, or in other words, decide whether there is exactly one optimal repair. The problem of repairing uniqueness (in a different repair semantics) is also referred to as *determinism* by Fan et al. [18]. In this paper, we study the three variants of this computational problem, under the three optimality semantics Pareto, global and completion, and denote them as *p-categoricity*, *g-categoricity* and *c-categoricity*, respectively.

It is known that under each of the three semantics there is always at least one preferred repair, and Staworko et al. [37] present a polynomial-time algorithm for finding such a repair. (We recall this algorithm in Section 3.) Hence, the categoricity problem is that of deciding whether the output of this algorithm is the only possible preferred repair. As we explain next, it turns out that each of the three variants of the problem entails quite a unique picture of complexity.

For the problem of p-categoricity, we focus on integrity constraints that are FDs, and establish the following dichotomy in data complexity, assuming that $P \neq NP$. For a relational schema with a set Δ of FDs:

- If Δ associates (up to equivalence) a single FD with every relation symbol, then p-categoricity is solvable

in polynomial time.

- In any other case, p-categoricity is coNP-complete.

For example, with the relation symbol $R(A, B, C)$ and the FD $A \rightarrow B$, p-categoricity is solvable in polynomial time; but if we add the dependency $B \rightarrow A$ then it becomes coNP-complete. Our proof uses a reduction technique from past dichotomies that involve FDs [12, 25], but requires some highly nontrivial additions.

We then turn to investigating c-categoricity, and establish a far more positive picture than the one for p-categoricity. In particular, the problem is solvable in polynomial time for every set of FDs. In fact, we present an algorithm for solving c-categoricity in polynomial time, assuming that constraints are given as an input *conflict hypergraph* [10]. (In particular, we establish polynomial-time data complexity for other types of integrity constraints, such as *conditional FDs* [5] and *denial constraints* [19].) The algorithm is extremely simple, yet its proof of correctness is quite intricate.

Finally, we explore g-categoricity, and focus first on FDs. We show that in the tractable case of p-categoricity (equivalence to a single FD per relation), g-categoricity is likewise solvable in polynomial time. For example, $R(A, B, C)$ with the dependency $A \rightarrow B$ has polynomial-time g-categoricity. Nevertheless, we prove that if we add the dependency $\emptyset \rightarrow C$ (that is, the attribute C should have the same value across all tuples), then g-categoricity becomes Π_2^P -complete. We do not complete a dichotomy as in p-categoricity, and leave that open for future work. Lastly, we observe that in our proof of Π_2^P -hardness, our reduction constructs a non-transitive priority relation, and we ask whether transitivity makes a difference. The three semantics of repairs remain different in the presence of transitivity. In particular, we show such a case where there are globally-optimal repairs that are not completion optimal repairs. Nevertheless, quite interestingly, we are able to prove that g-categoricity and c-categoricity are actually *the same problem* if transitivity is assumed. In particular, we establish that in the presence of transitivity, g-categoricity is solvable in polynomial time, even when constraints are given as a conflict hypergraph.

2. PRELIMINARIES

We now present some general terminology and notation that we use throughout the paper.

2.1 Signatures and Instances

A (*relational*) *signature* is a finite set $\mathcal{R} = \{R_1, \dots, R_n\}$ of *relation symbols*, each with a designated positive integer as its *arity*, denoted $\text{arity}(R_i)$. We assume an infinite set Const of *constants*, used as database values. An *instance* I over a signature $\mathcal{R} = \{R_1, \dots, R_n\}$ consists of finite relations $R_i^I \subseteq \text{Const}^{\text{arity}(R_i)}$, where $R_i \in \mathcal{R}$. We write $\llbracket R_i \rrbracket$ to denote the set $\{1, \dots, \text{arity}(R_i)\}$, and we refer to the members of $\llbracket R_i \rrbracket$ as *attributes* of R_i . If I is an instance over \mathcal{R} and \mathbf{t} is a tuple in R_i^I , then we say that $R_i(\mathbf{t})$ is a *fact* of I . By a slight abuse of notation, we identify an instance I with the set of its facts. For example, $R_i(\mathbf{t}) \in I$ denotes that

$R_i(t)$ is a fact of I . As another example, $J \subseteq I$ means that $R_i^J \subseteq R_i^I$ for every $R_i \in \mathcal{R}$; in this case, we say that J is *subinstance* of I .

In our examples, we often name the attributes and refer to them by their names. For instance, in Figure 1 we refer to the relation symbol as *CompCEO*(company, ceo) where company and ceo refer to Attributes 1 and 2, respectively. In the case of generic relation symbols, we implicitly name their attributes by capital English letters with the corresponding numeric values; for instance, we may refer to Attributes 1, 2 and 3 of $R/3$ by A , B and C , respectively. We stress that attribute names are not part of our formal model, but are rather used for readability.

2.2 Integrity and Inconsistency

Let \mathcal{R} be a signature, and I an instance over \mathcal{R} . In this paper we consider two representation systems for integrity constraints. The first is *functional dependencies* and the second is *conflict hypergraphs*.

2.2.1 Functional Dependencies

Let \mathcal{R} be a signature. A *Functional Dependency* (FD for short) over \mathcal{R} is an expression of the form $R : X \rightarrow Y$, where R is a relation symbol of \mathcal{R} , and X and Y are subsets of $\llbracket R \rrbracket$. When R is clear from the context, we may omit it and write simply $X \rightarrow Y$. A special case of an FD is a *key constraint*, which is an FD of the form $R : X \rightarrow Y$ where $X \cup Y = \llbracket R \rrbracket$. An FD $R : X \rightarrow Y$ is *trivial* if $Y \subseteq X$; otherwise, it is *nontrivial*.

When we are using the alphabetic attribute notation, we may write X and Y by simply concatenating the attribute symbols. For example, if we have a relation symbol $R/3$, then $A \rightarrow BC$ denotes the FD $R : \{1\} \rightarrow \{2, 3\}$. An instance I over R *satisfies* an FD $R : X \rightarrow Y$ if for every two facts f and g over R , if f and g agree on (i.e., have the same values for) the attributes of X , then they also agree on the attributes of Y . We say that I satisfies a set Δ of FDs if I satisfies every FD in Δ ; otherwise, we say that I *violates* Δ . Two sets Δ and Δ' of FDs are *equivalent* if for every instance I over \mathcal{R} it holds that I satisfies Δ if and only if it satisfies Δ' . For example, for $R/3$ the sets $\{A \rightarrow BC, C \rightarrow A\}$ and $\{A \rightarrow C, C \rightarrow AB\}$ are equivalent.

In this work, a *schema* \mathbf{S} is a pair (\mathcal{R}, Δ) , where \mathcal{R} is a signature and Δ is a set of FDs over \mathcal{R} . If $\mathbf{S} = (\mathcal{R}, \Delta)$ is a schema and $R \in \mathcal{R}$, then we denote by $\Delta|_R$ the restriction of Δ to the FDs $R : X \rightarrow Y$ over R .

EXAMPLE 2.1. In our first running example, we use the schema $\mathbf{S} = (\mathcal{R}, \Delta)$, defined as follows. The signature \mathcal{R} consists of a single relation *CompCEO*(company, ceo), which associates companies with their Chief Executive Officers (CEO). Figure 1 depicts an instance I over \mathcal{R} . We define Δ as the following set of FDs over \mathcal{R} .

$$\Delta \stackrel{\text{def}}{=} \{\text{company} \rightarrow \text{ceo}, \text{ceo} \rightarrow \text{company}\}$$

Hence, Δ states that in *CompCEO*, each company has a single CEO and each CEO manages a single company. Observe that I violates Δ . For example, Google has three CEOs,

<i>CompCEO</i>		
	company	ceo
f_{pi}^g	Google	Pichai
f_{pa}^g	Google	Page
f_{br}^g	Google	Brin
f_{pa}^a	Alphabet	Page
f_{pi}^a	Alphabet	Pichai

Figure 1: Inconsistent database of the company-CEO running example

Alphabet has two CEOs, and each of Pichai and Page is the CEO of two companies. \square

2.2.2 Conflict Hypergraphs

While FDs define integrity logically, at the level of the signature, a *conflict hypergraph* [10] provides a direct specification of inconsistencies at the instance level, by explicitly stating sets of tuples that cannot co-exist. In the case of FDs, the conflict hypergraph is a graph that has an edge between every two facts that violate an FD. Formally, for an instance I over a signature \mathcal{R} , a *conflict hypergraph* \mathcal{H} (for I) is a hypergraph that has the facts of I as its node set. A subinstance J of I is *consistent* with respect to (w.r.t.) \mathcal{H} if J is an *independent set* of \mathcal{H} ; that is, no hyperedge of \mathcal{H} is a subset of J . We say that J is *maximal* if $J \cup \{f\}$ is inconsistent for every $f \in I \setminus J$. When all the edges of a conflict hypergraph are of size two, we may call it a *conflict graph*.

Recall that conflict hypergraphs can represent inconsistencies for various types of integrity constraints, including FDs, the more general *conditional FDs* [5], and the more general *denial constraints* [19]. In fact, every constraint that is anti-monotonic (i.e., where subsets of consistent sets are always consistent) can be represented as a conflict hypergraph. In the case of denial constraints, the translation from the logical constraints to the conflict hypergraph can be done in polynomial time under *data complexity* (i.e., when the signature and constraints are assumed to be fixed).

Let $\mathbf{S} = (\mathcal{R}, \Delta)$ be a schema, and let I be an instance over \mathbf{S} . Recall that \mathbf{S} is assumed to have only FDs. We denote by $\mathcal{H}_{\mathbf{S}}^I$ the conflict graph for I that has an edge between every two facts that violate some FD of \mathbf{S} . Note that a subinstance J of I satisfies Δ if and only if J is consistent w.r.t. $\mathcal{H}_{\mathbf{S}}^I$. As an example, the left graph of Figure 2 depicts the graph $\mathcal{H}_{\mathbf{S}}^I$ for our running example; for now, the reader should ignore the directions on the edges, and view the graph as an undirected one. The following example involves a conflict hypergraph that is not a graph.

EXAMPLE 2.2. In our second running example, we use the toy scenario where the signature has a single relation symbol *Follows*/2, where *Follows*(x, y) means that person x follows person y (e.g., in a social network). We have two sets of people: a_i for $i = 1, 2, 3$, and b_j for $j = 1, \dots, 5$. All the facts have the form *Follows*(a_i, b_j); we denote such a fact by f_{ij} . The instance I has the following facts:

$$f_{11}, f_{12}, f_{21}, f_{22}, f_{23}, f_{24}, f_{31}, f_{32}, f_{34}, f_{35}$$

The hypergraph \mathcal{H} for I encodes the following rules:

- Each a_i can follow at most i people.
- Each b_j can be followed by at most j people.

Specifically, \mathcal{H} contains the following hyperedges:

- $\{f_{11}, f_{12}\}, \{f_{21}, f_{22}, f_{23}\}, \{f_{21}, f_{22}, f_{24}\},$
 $\{f_{21}, f_{23}, f_{24}\}, \{f_{22}, f_{23}, f_{24}\}, \{f_{31}, f_{32}, f_{34}, f_{35}\}$
- $\{f_{11}, f_{21}\}, \{f_{11}, f_{31}\}, \{f_{21}, f_{31}\}, \{f_{12}, f_{22}, f_{32}\}$

An example of a consistent subinstance J is

$$\{f_{11}, f_{22}, f_{23}, f_{32}, f_{34}, f_{35}\}.$$

The reader can verify that J is maximal. \square

2.3 Prioritizing Inconsistent Databases

We now recall the framework of preferred repairs by Staworko et al. [37]. Let I be an instance over a signature \mathcal{R} . A *priority relation* \succ over I is an acyclic binary relation over the facts in I . By *acyclic* we mean that I does not contain any sequence f_1, \dots, f_k of facts such that $f_i \succ f_{i+1}$ for all $i = 1, \dots, k-1$ and $f_k \succ f_1$. If \succ is a priority relation over I and K is a subinstance of I , then $\max_{\succ}(K)$ denotes the set of tuples $f \in K$ such that no $g \in K$ satisfies $g \succ f$.

An *inconsistent prioritizing instance* over \mathcal{R} is a triple (I, \mathcal{H}, \succ) , where I is an instance over \mathcal{R} , \mathcal{H} is a conflict hypergraph for I , and \succ a priority relation over I with the following property: for every two facts f and g in I , if $f \succ g$ then f and g are neighbors in \mathcal{H} (that is, f and g co-occur in some hyperedge).² For example, if $\mathcal{H} = \mathcal{H}_S^I$ (where all the constraints in S are FDs), then $f \succ g$ implies that $\{f, g\}$ violates at least one FD.

EXAMPLE 2.3. We continue our running company-CEO example. We define a priority relation \succ by $f_{pi}^g \succ f_{pa}^g$, $f_{pa}^g \succ f_{br}^g$ and $f_{pa}^a \succ f_{pi}^a$. We denote \succ by corresponding arrows on the left graph of Figure 2. (Therefore, some of the edges are directed and some are undirected.) We then get the inconsistent prioritizing instance $(I, \mathcal{H}_S^I, \succ)$ over \mathcal{R} . Observe that the graph does not contain directed cycles, as required from a priority relation. \square

EXAMPLE 2.4. Recall that the instance I of our followers example is defined in Example 2.2. The priority relation \succ is given by $f_{il} \succ f_{jk}$ if one of the following holds: (a) $i = j$ and $k = l + 1$, or (b) $j = i + 1$ and $l = k$. For example, we have $f_{11} \succ f_{12}$ and $f_{12} \succ f_{22}$. But we do not have $f_{11} \succ f_{22}$ (hence, \succ is not transitive). \square

Let (I, \mathcal{H}, \succ) be an inconsistent prioritizing instance over a signature \mathcal{R} . We say that \succ is *total* if for every two facts f and g in I , if f and g are neighbors then either $f \succ g$ or $g \succ f$. A priority \succ_c over I is a *completion* of \succ (w.r.t. \mathcal{H}) if \succ is a subset of \succ_c and \succ_c is total. As an example, the middle and right graphs of Figure 2 are two completions of the

²This requirement has been made with the introduction of the framework [37]. Obviously, the lower bounds we present hold even without this requirement. Moreover, our main upper bound, Theorem 6.1, holds as well without this requirement. We defer to future work the thorough investigation of the impact of relaxing this requirement.




Figure 2: The conflict graph \mathcal{H}_S^I and the priority relation \succ for the company-CEO running example (left), and two completions of \succ (middle and right)

priority relation \succ depicted on the left side. A *completion* of (I, \mathcal{H}, \succ) is an inconsistent prioritizing instance $(I, \mathcal{H}, \succ_c)$ where \succ_c is a completion of \succ .

2.4 Preferred Repairs

Let $D = (I, \mathcal{H}, \succ)$ be an inconsistent prioritizing instance over \mathcal{R} . As defined by Arenas et al. [3], J is a *repair* of D if J is a maximal consistent subinstance of I . Staworko et al. [37] define three different notions of *preferred repairs*: *Pareto optimal*, *globally optimal*, and *completion optimal*. The first two notions are based on checking whether a repair J of I can be improved by replacing a set of facts in J with a more preferred set of facts from I . They differ by the way they define when one set of facts is considered more preferred than another one. The last notion is based on the notion of completion. Next we give the formal definitions.

DEFINITION 2.5 (IMPROVEMENT). Let (I, \mathcal{H}, \succ) be an inconsistent prioritizing instance over a signature \mathcal{R} , and J and J' two distinct consistent subinstances of I .

- J is a *Pareto improvement* of J' if there exists a fact $f \in J \setminus J'$ such that $f \succ f'$ for all facts $f' \in J' \setminus J$.
- J is a *global improvement* of J' if for every fact $f' \in J' \setminus J$ there exists a fact $f \in J \setminus J'$ such that $f \succ f'$.

That is, J is a Pareto improvement of J' if, in order to obtain J from J' , we insert and delete facts, and one of the inserted facts is preferred to all deleted facts. And J is a global improvement of J' if, in order to obtain J from J' , we insert and delete facts, and every deleted fact is preferred to by some inserted fact.

EXAMPLE 2.6. We continue the company-CEO running example. We define three consistent subinstances of I .

$$J_1 \stackrel{\text{def}}{=} \{f_{br}^g, f_{pi}^a\} \quad J_2 \stackrel{\text{def}}{=} \{f_{pa}^g, f_{pi}^a\}$$

$$J_3 \stackrel{\text{def}}{=} \{f_{br}^g, f_{pa}^a\} \quad J_4 \stackrel{\text{def}}{=} \{f_{pi}^g, f_{pa}^a\}$$

Note the following. First, J_2 is a Pareto improvement of J_1 , since $f_{pa}^g \in J_2 \setminus J_1$ and $f_{pa}^g \succ f$ for every fact in $J_1 \setminus J_2$ (where in this case there is only one such an f , namely f_{br}^g). Second, J_4 is a global improvement of J_2 because $f_{pi}^g \succ f_{pa}^g$ and $f_{pa}^a \succ f_{pi}^a$. (We refer to J_3 in later examples.) \square

We then get the following variants of *preferred repairs*.

DEFINITION 2.7 (P/G/C-REPAIR). Let D be an inconsistent prioritizing instance (I, \mathcal{H}, \succ) , and let J be a consistent subinstance of I . Then J is a:

- Pareto-optimal repair of D if there is no Pareto improvement of J .
- globally-optimal repair of D if there is no global improvement of J .
- completion-optimal repair of D if there exists a completion D_c of D such that J is a globally-optimal repair of D_c .

We abbreviate “Pareto-optimal repair,” “globally-optimal repair,” and “completion-optimal repair” by p-repair, g-repair and c-repair, respectively.

We remark that in the definition of a completion-optimal repair, we could replace “globally-optimal” with “Pareto-optimal” and obtain an equivalent definition [37].

Let $D = (I, \mathcal{H}, \succ)$ be an inconsistent prioritizing instance over a signature \mathcal{R} . We denote the set of all the repairs, p-repairs, g-repairs and c-repairs of D by $\text{Rep}(D)$, $\text{PRep}(D)$, $\text{GRep}(D)$ and $\text{CRep}(D)$, respectively. The following was shown by Staworko et al. [37].

PROPOSITION 2.8. [37] For all inconsistent prioritizing instances D we have $\text{CRep}(D) \neq \emptyset$, and moreover,

$$\text{CRep}(D) \subseteq \text{GRep}(D) \subseteq \text{PRep}(D) \subseteq \text{Rep}(D).$$

EXAMPLE 2.9. We continue our company-CEO example. Recall the instances J_i defined in Example 2.6. We have shown that J_1 has a Pareto improvement, and therefore, J_1 is not a p-repair (although it is a repair in the ordinary sense). The reader can verify that J_2 has no Pareto improvements, and therefore, it is a p-repair. But J_2 is not a g-repair, since J_4 is a global improvement of J_2 . The reader can verify that J_3 is a g-repair (hence, a p-repair). Finally, observe that J_4 is a g-repair w.r.t. the left completion of \succ in Figure 2 (and also w.r.t. the right one). Hence, J_4 is a c-repair (hence, a g-repair and a p-repair). In contrast, observe that J_3 has a global improvement (and a Pareto improvement) in both completions; but it does not prove that J_3 is not a c-repair (since, conceptually, one needs to consider all possible completions of \succ). \square

EXAMPLE 2.10. We now continue the follower example. The inconsistent prioritizing instance (I, \mathcal{H}, \succ) is defined in Examples 2.2 and 2.4. Consider the following instance.

$$J_1 \stackrel{\text{def}}{=} \{f_{11}, f_{22}, f_{23}, f_{32}, f_{34}, f_{35}\}$$

The reader can verify that J_1 is a c-repair (e.g., by completing \succ through the lexicographic order). The subinstance $J_2 = \{f_{12}, f_{21}, f_{22}, f_{34}, f_{35}\}$ is a repair but not a p-repair, since we can add f_{11} and remove both f_{12} and f_{21} , and thus obtain a Pareto improvement. \square

3. CATEGORICITY

In this section we define the computational problem of *categoricity*, which is the main problem that we study in this paper. Proposition 2.8 states that, under each of the semantics of preferred repairs, at least one such a repair exists. In general, there can be many possible preferred repairs. The problem of *categoricity* [37] is that of testing whether there is *precisely* one such a repair; that is, there do not exist two distinct preferred repairs, and therefore, the priority relation contains enough information to clean the inconsistent instance unambiguously.

PROBLEM 3.1. *Problems p-categoricity, g-categoricity, and c-categoricity are those of testing whether $|\text{PRep}(D)| = 1$, $|\text{GRep}(D)| = 1$ and $|\text{CRep}(D)| = 1$, respectively, given a signature \mathcal{R} and an inconsistent prioritizing instance D over \mathcal{R} .*

As defined, categoricity takes as input both the signature \mathcal{R} and the inconsistent prioritizing instance D , where constraints are represented by a conflict hypergraph. We also study this problem from the perspective of *data complexity*, where we fix a schema $\mathbf{S} = (\mathcal{R}, \Delta)$, where Δ is a set of FDs. In that case, the input consists of an instance I over \mathcal{R} and a priority relation \prec over I . The conflict hypergraph is then implicitly assumed to be $\mathcal{H}_{\mathbf{S}}^I$. We denote the corresponding variants of the problem by p-categoricity $\langle \mathbf{S} \rangle$, g-categoricity $\langle \mathbf{S} \rangle$ and c-categoricity $\langle \mathbf{S} \rangle$, respectively.

EXAMPLE 3.2. Continuing our company-CEO example, we showed in Example 2.9 that there are at least two g-repairs and at least three p-repairs. Hence, a solver for g-categoricity $\langle \mathbf{S} \rangle$ should return false on (I, \succ) , and so is a solver for p-categoricity $\langle \mathbf{S} \rangle$. In contrast, we will later show that there is precisely one c-repair (Example 6.2); hence, a solver for c-categoricity $\langle \mathbf{S} \rangle$ should return true on (I, \succ) . If, on the other hand, we replaced \succ with any of the completions in Figure 2, then there would be precisely one p-repair and one g-repair (namely, the current single c-repair). This follows from a result of Staworko et al. [37], stating that categoricity holds in the case of total priority relations. \square

4. PRELIMINARY INSIGHTS

We begin with some basic insights into the different variants of the categoricity problem.

4.1 Greedy Repair Generation

We recall an algorithm by Staworko et al. [37] for greedily constructing a c-repair. This is the algorithm FindCRep of Figure 3. The algorithm takes as input an inconsistent prioritizing instance (I, \mathcal{H}, \succ) and returns a c-repair J . It begins with an empty J , and incrementally inserts tuples to J , as follows. In each iteration of lines 3–6, the algorithm selects a fact f from $\max_{\succ}(I)$ and removes it from I . Then, f is added to J if it does not violate consistency, that is, if \mathcal{H} does not contain any hyperedge e such that $e \subseteq J \cup \{f\}$. The specific way of choosing the fact f among all those in $\max_{\succ}(I)$ is (deliberately) left unspecified, and hence, different executions may result in different c-repairs. In that

Algorithm FindCRep(I, \mathcal{H}, \succ)

```

1:  $J := \emptyset$ 
2: while  $\max_{\succ}(I) \neq \emptyset$  do
3:   choose a fact  $f$  in  $\max_{\succ}(I)$ 
4:    $I := I \setminus \{f\}$ 
5:   if  $J \cup \{f\}$  is consistent w.r.t.  $\mathcal{H}$  then
6:      $J := J \cup \{f\}$ 
7: return  $J$ 

```

Figure 3: Finding a c-repair [37]

sense, the algorithm is nondeterministic. Staworko et al. [37] proved that the possible results of these different executions are *precisely* the c-repairs.

THEOREM 4.1. [37] *Let (I, \mathcal{H}, \succ) be an inconsistent prioritizing instance over \mathcal{R} . Let J be a consistent subinstance of I . Then J is a c-repair if and only if there exists an execution of FindCRep(I, \mathcal{H}, \succ) that returns J .*

Due to Theorem 4.1, we often refer to a c-repair as a *greedy* repair. This theorem, combined with Proposition 2.8, has several implications for us. First, we can obtain an x-repair (where x is either p, g or c) in polynomial time. Hence, if a solver for x-categoricity determines that there is a single x-repair, then we can actually generate that x-repair in polynomial time. Second, c-categoricity is the problem of testing whether FindCRep(I, \mathcal{H}, \succ) returns the same instance J on every execution. Moreover, due to Proposition 2.8, p-categoricity (resp. g-categoricity) is the problem of testing whether every p-repair (resp. g-repair) is equal to the one that is obtained by some execution of the algorithm.

EXAMPLE 4.2. We consider the application of the algorithm FindCRep to the instance of our company-CEO example (where $\mathcal{H} = \mathcal{H}_S^I$). The following are two different executions. We denote inclusion in J (i.e., the condition of line 5 is true) by plus and exclusion from J by minus.

- $+f_{pi}^g, -f_{pa}^g, -f_{br}^g, +f_{pa}^a, -f_{pi}^a$.
- $+f_{pa}^a, -f_{pi}^a, +f_{pi}^g, -f_{pa}^g, -f_{br}^g$.

Observe that both executions return $J_4 = \{f_{pi}^g, f_{pa}^a\}$. This is in par with the statement in Example 3.2 that in this running example there is a single c-repair. \square

4.2 Complexity Insights

Our goal is to study the complexity of x-categoricity (where x is g, p and c). This problem is related to that of *x-repair checking*, namely, given $D = (I, \mathcal{H}, \succ)$ and J , determine whether J is an x-repair of D . The following is known about this problem.

THEOREM 4.3. [12, 37] *The following hold.*

- *p-repair checking and c-repair checking are solvable in polynomial time; g-repair checking is in coNP [37].*

- *Let $\mathbf{S} = (\mathcal{R}, \Delta)$ be a fixed schema. If $\Delta_{|R}$ is equivalent to either a single FD or two key constraints for every $R \in \mathcal{R}$, then g-repair checking is solvable in polynomial time; otherwise, g-repair checking is coNP-complete [12].*

Recall from Proposition 2.8 that there is always at least one x-repair. Therefore, given (I, \mathcal{H}, \succ) we can solve the problem using a coNP algorithm with an oracle to x-repair checking: for all two distinct subinstances J_1 and J_2 , either J_1 or J_2 is not an x-repair. Therefore, from Theorem 4.3 we conclude the following.

COROLLARY 4.4. *The following hold.*

- *p-categoricity and c-categoricity are in coNP.*
- *For all fixed schemas $\mathbf{S} = (\mathcal{R}, \Delta)$, g-categoricity(\mathbf{S}) is in Π_2^p , and moreover, if $\Delta_{|R}$ is equivalent to either a single FD or two key constraints for every $R \in \mathcal{R}$ then g-categoricity(\mathbf{S}) is in coNP.*

We stress here that if x-categoricity is solvable in polynomial time, then x-categoricity(\mathbf{S}) is solvable in polynomial time for *all* schemas \mathbf{S} ; this is true since for every fixed schema \mathbf{S} the hypergraph \mathcal{H}_S^I can be constructed in polynomial time, given I . Similarly, if x-categoricity(\mathbf{S}) is coNP-hard (resp. Π_2^p -hard) for *at least one* \mathbf{S} , then x-categoricity is coNP-hard (resp. Π_2^p -hard).

When we are considering x-categoricity(\mathbf{S}), we assume that all the integrity constraints are FDs. Therefore, unlike the general problem of x-categoricity, in x-categoricity(\mathbf{S}) conflicting facts always belong to the same relation. It thus follows that our analysis for x-categoricity(\mathbf{S}) can restrict to single-relation schemas. Formally, we have the following.

PROPOSITION 4.5. *Let $\mathbf{S} = (\mathcal{R}, \Delta)$ be a schema and x be one of p, g and c. For each relation $R \in \mathcal{R}$, let $\mathbf{S}_{|R}$ be the schema $(\{R\}, \Delta_{|R})$.*

- *If x-categoricity($\mathbf{S}_{|R}$) is solvable in polynomial time for every $R \in \mathcal{R}$, then x-categoricity(\mathbf{S}) is solvable in polynomial time.*
- *If x-categoricity($\mathbf{S}_{|R}$) is coNP-hard (resp. Π_2^p -hard) for at least one $R \in \mathcal{R}$, then x-categoricity(\mathbf{S}) is coNP-hard (resp. Π_2^p -hard).*

Observe that the phenomenon of Proposition 4.5 *does not* hold for x-categoricity, since the given conflict hypergraph may include hyperedges that cross relations.

In the following sections we investigate each of the three variants of categoricity: p-categoricity (Section 5), c-categoricity (Section 6) and g-categoricity (Section 7).

5. P-CATEGORICITY

In this section we prove a dichotomy in the complexity of p-categoricity(\mathbf{S}) over all schemas \mathbf{S} (where Δ consists of FDs). This dichotomy states that the only tractable case is where the schema associates a single FD (which can be trivial) to each relation symbol, up to equivalence. In all other cases, p-categoricity(\mathbf{S}) is coNP-complete. Formally, we prove the following.

THEOREM 5.1. *Let $\mathbf{S} = (\mathcal{R}, \Delta)$ be a schema. The problem $p\text{-categoricity}(\mathbf{S})$ can be solved in polynomial time if $\Delta|_R$ is equivalent to a single FD for every $R \in \mathcal{R}$. In every other case, $p\text{-categoricity}(\mathbf{S})$ is coNP-complete.*

The proof of Theorem 5.1 is involved, and we outline it in the rest of this section. The tractability side is fairly simple (as we show in the next section), and the challenge is in the hardness side. Due to Proposition 4.5, it suffices to consider schemas \mathbf{S} with a single relation. Hence, in the remainder of this section we consider only such schemas \mathbf{S} .

5.1 Proof of Tractability

In this section we fix a schema $\mathbf{S} = (\mathcal{R}, \Delta)$, such that \mathcal{R} consist of a single relational symbol R . We will prove that $p\text{-categoricity}(\mathbf{S})$ is solvable in polynomial time if Δ is a singleton. We denote the single FD in Δ as $X \rightarrow Y$. We fix the input (I, \succ) for $p\text{-categoricity}(\mathbf{S})$.

For a fact $f \in R^I$, we denote by $f[X]$ and $f[Y]$ the restriction of the tuple of f to the attributes in X and Y , respectively. Adopting the terminology of Koutris and Wijsen [29], a *block* of I is a maximal collection of facts of I that agree on all the attributes of X (i.e., facts f that have the same $f[X]$). Similarly, a *subblock* of I is a maximal collection of facts that agree on both X and Y . For tuples \mathbf{a} and \mathbf{b} constants, we denote by $I_{\mathbf{a}}$ the block of facts f with $f[X] = \mathbf{a}$, and by $I_{\mathbf{a}, \mathbf{b}}$ the subblock of facts f with $f[X] = \mathbf{a}$ and $f[Y] = \mathbf{b}$.

EXAMPLE 5.2. Consider again the instance I of Figure 1, and suppose that Δ consists of only $\text{company} \rightarrow \text{ceo}$ (i.e., each company has a single CEO, but a person can be the CEO of several companies). Then for $\mathbf{a} = (\text{Google})$ and $\mathbf{b} = (\text{Pichai})$ the block $I_{\mathbf{a}}$ is $\{f_{\text{pi}}^g, f_{\text{pa}}^g, f_{\text{br}}^g\}$ and the subblock $I_{\mathbf{a}, \mathbf{b}}$ is the singleton $\{f_{\text{pi}}^g\}$. \square

Tractability for \mathbf{S} is based on the following lemma.

LEMMA 5.3. *Let J be a subinstance of I . Then J is a $p\text{-repair}$ if and only if J is a union of $p\text{-repairs}$ over all the blocks $I_{\mathbf{a}}$ of I . Moreover, each $p\text{-repair}$ of a block $I_{\mathbf{a}}$ is a subblock $I_{\mathbf{a}, \mathbf{b}}$.*

We then get the following lemma.

LEMMA 5.4. *The following are equivalent.*

1. I has a single $p\text{-repair}$.
2. Each block $I_{\mathbf{a}}$ has a single $p\text{-repair}$.
3. No block $I_{\mathbf{a}}$ has two distinct subblocks $I_{\mathbf{a}, \mathbf{b}}$ and $I_{\mathbf{a}, \mathbf{c}}$ that are $p\text{-repairs}$ of $I_{\mathbf{a}}$.

A polynomial-time algorithm then follows directly from Lemma 5.4 and the fact that $p\text{-repair}$ checking is solvable in polynomial time (Theorem 4.3).

5.2 Proof of Hardness

The hardness side of the dichotomy is more involved than its tractability side. Our proof is based on the concept of a *fact-wise reduction* [26], which has also been used by Fagin et al. [12] in the context of $g\text{-repair}$ checking.

5.2.1 Fact-Wise Reduction

Let $\mathbf{S} = (\mathcal{R}, \Delta)$ and $\mathbf{S}' = (\mathcal{R}', \Delta')$ be two schemas. A *mapping* from \mathcal{R} to \mathcal{R}' is a function μ that maps facts over \mathcal{R} to facts over \mathcal{R}' . We naturally extend a mapping μ to map instances J over \mathcal{R} to instances over \mathcal{R}' by defining $\mu(J)$ to be $\{\mu(f) \mid f \in J\}$. A *fact-wise reduction* from \mathbf{S} to \mathbf{S}' is a mapping Π from \mathcal{R} to \mathcal{R}' with the following properties.

1. Π is injective; that is, for all facts f and g over \mathcal{R} , if $\Pi(f) = \Pi(g)$ then $f = g$.
2. Π preserves consistency and inconsistency; that is, for every instance J over \mathbf{S} , the instance $\Pi(J)$ satisfies Δ' if and only if J satisfies Δ .
3. Π is computable in polynomial time.

Let \mathbf{S} and \mathbf{S}' be two schemas, and let Π be a fact-wise reduction from \mathbf{S} to \mathbf{S}' . Given an inconsistent instance I over \mathbf{S} and a priority relation \succ over I , we denote by $\Pi(\succ)$ the priority relation \succ' over $\Pi(I)$ where $\Pi(f) \succ' \Pi(g)$ if and only if $f \succ g$. If D is the inconsistent prioritizing instance $(I, \mathcal{H}_{\mathbf{S}}^I, \succ)$, then we denote by $\Pi(D)$ the triple $(\Pi(I), \mathcal{H}_{\mathbf{S}'}^{\Pi(I)}, \Pi(\succ))$, which is also an inconsistent prioritizing instance. The usefulness of fact-wise reductions is due to the following proposition, which is straightforward.

PROPOSITION 5.5. *Let \mathbf{S} and \mathbf{S}' be two schemas, and suppose that Π is a fact-wise reduction from \mathbf{S} to \mathbf{S}' . Let I be an inconsistent instance over \mathbf{S} , \succ a priority relation over I , and D the inconsistent prioritizing instance $(I, \mathcal{H}_{\mathbf{S}}^I, \succ)$. Then there is a bijection between $\text{PRep}(D)$ and $\text{PRep}(\Pi(D))$.*

We then conclude the following corollary.

COROLLARY 5.6. *If there is a fact-wise reduction from \mathbf{S} to \mathbf{S}' , then there is a polynomial-time reduction from $p\text{-categoricity}(\mathbf{S})$ to $p\text{-categoricity}(\mathbf{S}')$.*

5.2.2 Specific Schemas

In the proof we consider seven specific schemas. The importance of these schemas will later become apparent. We denote these schemas by \mathbf{S}^i , for $i = 0, 1, \dots, 6$, where each \mathbf{S}^i is the schema $(\mathcal{R}^i, \Delta^i)$, and \mathcal{R}^i is the singleton $\{R^i\}$. The specification of the \mathbf{S}^i is as follows.

0. $R^0/2$ and $\Delta^0 = \{A \rightarrow B, B \rightarrow A\}$
1. $R^1/3$ and $\Delta^1 = \{AB \rightarrow C, BC \rightarrow A, AC \rightarrow B\}$
2. $R^2/3$ and $\Delta^2 = \{A \rightarrow B, B \rightarrow A\}$
3. $R^3/3$ and $\Delta^3 = \{AB \rightarrow C, C \rightarrow B\}$
4. $R^4/3$ and $\Delta^4 = \{A \rightarrow B, B \rightarrow C\}$
5. $R^5/3$ and $\Delta^5 = \{A \rightarrow C, B \rightarrow C\}$
6. $R^6/3$ and $\Delta^6 = \{\emptyset \rightarrow A, B \rightarrow C\}$

(In the definition of \mathbf{S}^6 , recall that $\emptyset \rightarrow A$ denotes the FD $\emptyset \rightarrow \{1\}$, meaning that all tuples should have the same value for their first attribute.) In the proof we use fact-wise reductions from the \mathbf{S}^i , as we explain in the next section.




Figure 4: The structure of fact-wise reductions for proving the hardness side of the dichotomy of Theorem 5.1

5.2.3 Two Hard Schemas

Our proof boils down to proving coNP-hardness for two specific schemas, namely S^0 and S^6 , and then using (known and new) fact-wise reductions in order to cover all the other schemas. For S^6 the proof is fairly simple. But hardness for S^0 turns out to be quite challenging to prove, and in fact, this part is the hardest in the proof of Theorem 5.1. Note that S^0 is the schema of our company-CEO running example (introduced in Example 2.1).

THEOREM 5.7. *The problems $p\text{-categoricity}(S^0)$ and $p\text{-categoricity}(S^6)$ are both coNP-hard.*

The proof (as well as all the other proofs for the results in this paper) can be found in the appendix.

5.2.4 Applying Fact-Wise Reductions

The following has been proved by Fagin et al. [12].

THEOREM 5.8. [12] *Let $S = (\mathcal{R}, \Delta)$ be a schema such that \mathcal{R} consists of a single relation symbol. Suppose that Δ is equivalent to neither any single FD nor any pair of keys. Then there is a fact-wise reduction from some S^i to S , where $i \in \{1, \dots, 6\}$.*

In the appendix we prove the following two lemmas, giving additional fact-wise reductions.

LEMMA 5.9. *Let $S = (\mathcal{R}, \Delta)$ be a schema such that \mathcal{R} consists of a single relation symbol. Suppose that Δ is equivalent to a pair of keys, and Δ is not equivalent to any single FD. Then there is a fact-wise reduction from S^0 to S .*

LEMMA 5.10. *For all $i = 1, \dots, 5$ there is a fact-wise reduction from S^0 to S^i .*

The structure of our fact-wise reductions is depicted in Figure 4. Dashed edges are known fact-wise reductions,

Algorithm CCategory(I, \mathcal{H}, \succ)

```

1:  $i := 0$ 
2:  $J := \emptyset$ 
3: while  $I \neq \emptyset$  do
4:    $i := i + 1$ 
5:    $P_i := \max_{\succ^+}(I)$ 
6:    $J := J \cup P_i$ 
7:    $N_i := \{f \in I \mid \mathcal{H} \text{ has a hyperedge } e \text{ s.t. } f \in e, (e \setminus \{f\}) \subseteq J, \text{ and } (e \setminus \{f\}) \succ^+ f\}$ 
8:    $I := I \setminus (P_i \cup N_i)$ 
9: return true iff  $J$  is consistent

```

Figure 5: Algorithm for c-categoricity

while solid edges are novel. Observe that each single-relation schema on the hardness side of Theorem 5.1 has an ingoing path from either S^0 or S^6 , both shown to have coNP-hard p-categoricity (Theorem 5.7).

6. C-CATEGORICITY

We now investigate the complexity of c-categoricity. Our main result is that this problem is tractable.

THEOREM 6.1. *The c-categoricity problem is solvable in polynomial time.*

In the remainder of this section we establish Theorem 6.1 by presenting a polynomial-time algorithm for solving c-categoricity. The algorithm is very simple, but its proof of correctness (given in the appendix) is intricate.

6.1 Notation

To present our algorithm, some notation is required. Let (I, \mathcal{H}, \succ) be an inconsistent prioritizing instance. The *transitive closure* of \succ , denoted \succ^+ , is the priority relation over the facts of I where for every two facts f and g in I it holds that $f \succ^+ g$ if and only if there exists a sequence f_0, \dots, f_m of facts, where $m > 0$, such that $f = f_0$, $f_m = g$, and $f_i \succ f_{i+1}$ for all $i = 0, \dots, m-1$. Obviously, \succ^+ is acyclic (since \succ is acyclic). But unlike \succ , the relation \succ^+ may compare between facts that are not necessarily neighbors in \mathcal{H} .

Let (I, \mathcal{H}, \succ) be an inconsistent prioritizing instance, let K be a set of facts of I , and let f be a fact of I . By $K \succ^+ f$ we denote the fact that $g \succ^+ f$ for every fact $g \in K$.

6.2 Algorithm

Figure 5 depicts a polynomial-time algorithm for solving c-categoricity. We next explain how it works, and later discuss its correctness.

As required, the input for the algorithm is an inconsistent prioritizing instance (I, \mathcal{H}, \succ) . (The signature \mathcal{R} is not needed by the algorithm.) The algorithm incrementally constructs a subinstance J of I , starting with an empty J . Later we will prove that there is a single c-repair if and only if

J is consistent; and in that case, J is the single c-repair. The loop in the algorithm constructs fact sets P_1, \dots, P_t and N_1, \dots, N_t . Each P_i is called a *positive stratum* and each N_i is called a *negative stratum*. Both P_i and N_i are constructed in the i th iteration. On that iteration we add all the facts of P_i to J and remove from I all the facts of P_i and all the facts of N_i . The sets P_i and N_i are defined as follows.

- P_i consists of the maximal facts in the current I , according to \succ^+ .
- N_i consists of all the facts f that, together with $P_1 \cup \dots \cup P_i$, complete a hyperedge of preferred facts; that is, \mathcal{H} contains a hyperedge that contains f , is contained in $P_1 \cup \dots \cup P_i \cup \{f\}$, and satisfies $g \succ^+ f$ for every incident $g \neq f$.

The algorithm continues to iterate until I gets empty. As said above, in the end the algorithm returns true if J is consistent, and otherwise false. Next, we give some examples of executions of the algorithm.

EXAMPLE 6.2. Consider the inconsistent prioritizing instance (I, \mathcal{H}, \succ) from our company-CEO running example, illustrated on the left side of Figure 2. The algorithm makes a single iteration on this instance, where $P_1 = \{f_{pi}^g, f_{pa}^a\}$ and $N_1 = \{f_{pa}^g, f_{pi}^a, f_{pa}^g\}$. Both f_{pi}^g and f_{pa}^a are in P_1 since both are maximal. Also, each of f_{pa}^g, f_{pi}^a and f_{pa}^g is in conflict with P_1 , and we have $f_{pi}^g \succ f_{pa}^g, f_{pa}^a \succ f_{pi}^a$, and $f_{pi}^g \succ^+ f_{br}^g$. \square

EXAMPLE 6.3. Now consider the inconsistent prioritizing instance (I, \mathcal{H}, \succ) from our followers running example. Figure 6 illustrates the execution of the algorithm, where each column describes P_i or N_i , from left to right in the order of their construction. For convenience, the priority relation \succ , as defined in Example 2.4, is depicted in Figure 6 using corresponding edges between the facts.

On iteration 1, for instance, we have $P_1 = \{f_{11}, f_{34}\}$, since f_{11} and f_{34} are the facts without incoming edges on Figure 6. Moreover, we have $N_1 = \{f_{12}, f_{21}, f_{31}\}$. The reason why N_1 contains f_{12} , for example, is that $\{f_{11}, f_{12}\}$ is a hyperedge, the fact f_{11} is in P_1 , and $f_{11} \succ f_{12}$ (hence, $f_{11} \succ^+ f_{12}$). For a similar reason N_1 contains f_{21} . Fact f_{31} is in N_1 as $\{f_{11}, f_{31}\}$ is a hyperedge, and though $f_{11} \not\succ f_{31}$, we have $f_{11} \succ^+ f_{31}$. As another example, N_3 contains f_{24} since \mathcal{H} has the hyperedge $\{f_{22}, f_{23}, f_{24}\}$, the set $\{f_{22}, f_{23}\}$ is contained in $P_1 \cup P_2 \cup P_3$, and $\{f_{22}, f_{23}\} \succ^+ f_{24}$.

In the end, $J = \{f_{11}, f_{22}, f_{23}, f_{32}, f_{34}, f_{35}\}$, which is also the subinstance J_1 of Example 2.10. Since J is consistent, the algorithm will determine that there is a single c-repair, and that c-repair is J . \square

EXAMPLE 6.4. We now give an example of an execution on a negative instance of c-categoricity. (In Section 7 we refer to this example for a different reason.) Figure 7 shows an instance I over the schema S^1 , which is defined in Section 5.2.2. Recall that in this schema every two attributes form a key. Each fact $R^1(a_1, a_2, a_3)$ in I is depicted by a tuple that consists of the three values. For example, I contains the (conflicting) facts $R^1(A, a, 1)$ and $R^1(A, a, 2)$. Hereon, we write Xy_i instead of $R^1(X, y, i)$. The priority relation \succ is




Figure 6: Execution of CCategorycity on the followers example




Figure 7: An inconsistent instance I over S^1 with a priority relation \succ over I

given by the directed edges between the facts; for example, $Aa1 \succ Aa2$. Undirected edges are between conflicting facts that are incomparable by \succ (e.g., $Ab2$ and $Ab3$).

The execution of the algorithm on $(I, \mathcal{H}_{S^1}^I, \succ)$ is as follows. On the first iteration, $P_1 = \{Aa1, Ab2, Ba3, Bb1\}$ and $N_1 = \{Aa2, Bb3\}$. In particular, note that N_1 does not contain $Ba2$ since it conflicts only with $Ba3$ in P_1 , but the two are incomparable. Similarly, N_1 does not contain $Ab3$ since it is incomparable with $Ab2$. Consequently, in the second iteration we have $P_2 = \{Ba2, Ab3\}$ and $N_2 = \emptyset$. In the end, $J = P_1 \cup P_2$ is inconsistent, and therefore, the algorithm will return false. Indeed, the reader can easily verify that each of the following is a c-repair: $\{Aa1, Ab2, Ba3, Bb1\}$, $\{Aa1, Ab2, Ba2, Bb1\}$, and $\{Aa1, Ba3, Ab3, Bb1\}$. \square

6.3 Correctness

Correctness of CCategorycity is stated in the following theorem.

THEOREM 6.5. *Let (I, \mathcal{H}, \succ) be an inconsistent prioritizing instance, and let J be the subinstance of I constructed in the execution of CCategorycity(I, \mathcal{H}, \succ). Then the following are equivalent.*

1. J is consistent.
2. There is a single c-repair.

Moreover, if J is consistent then J is the single c-repair.

Theorem 6.5, combined with the observation that the algorithm CCategorycity terminates in polynomial time, imply Theorem 6.1. As previously said, the proof of Theorem 6.5 is quite involved. The direction $1 \rightarrow 2$ is that of *soundness*—if the algorithm returns true then there is precisely one c-

repair. The direction $2 \rightarrow 1$ is that of *completeness*—if there is precisely one c-repair then the algorithm returns true.

Soundness is the easier direction to prove. We assume, by way of contradiction, that there is a c-repair J' different from the subinstance J returned by the algorithm. Such J' must include a fact f' from some negative stratum. We consider an execution of the algorithm FindCRep that returns J' , and establish a contradiction by considering the first time such an f' is being added to the constructed solution.

Proving completeness is more involved. We assume, by way of contradiction, that the constructed J is inconsistent. We are looking at the first positive stratum P_i such that $P_1 \cup \dots \cup P_i$ contains a hyperedge. Then, the crux of the proof is in showing that we can then construct two c-repairs using the algorithm FindCRep: one contains some fact from P_i and another one does not contain that fact. We then establish that there are at least two c-repairs, hence a contradiction.

7. G-CATEGORICITY

In this section, we investigate the complexity of the g-categoricity. We first show a tractability result for the case of a schema with a single FD. Then, we show Π_2^P -completeness for a specific schema. Finally, we discuss the implication of assuming *transitivity* in the priority relation, and show a general positive result therein.

7.1 Tractable Schemas

Recall from Theorem 5.1 that, assuming $P \neq NP$, the problem p-categoricity $\langle S \rangle$ is solvable in polynomial time if and only if S consists (up to equivalence) of a single FD per relation. The reader can verify that the same proof works for g-categoricity $\langle S \rangle$. Hence, our first result is that the tractable schemas of p-categoricity remain tractable for g-categoricity.

THEOREM 7.1. *Let $S = (\mathcal{R}, \Delta)$ be a schema. The problem g-categoricity $\langle S \rangle$ can be solved in polynomial time if $\Delta|_R$ is equivalent to a single FD for every $R \in \mathcal{R}$.*

It is left open whether there is any schema S that is not as in Theorem 7.1 where g-categoricity $\langle S \rangle$ is solvable in polynomial time. In the next section we give an insight into this open problem (Theorem 7.4).

7.2 Intractable Schemas

Our next result shows that g-categoricity $\langle S \rangle$ hits a harder complexity class than p-categoricity $\langle S \rangle$. In particular, while p-categoricity $\langle S \rangle$ is always in coNP (due to Theorem 4.4), we will show a schema S where g-categoricity $\langle S \rangle$ is Π_2^P -complete. This schema is the schema S^6 from Section 5.2.2.

THEOREM 7.2. *g-categoricity $\langle S^6 \rangle$ is Π_2^P -complete.*

The proof of Theorem 7.2 is by a reduction from the Π_2^P -complete problem QCNF $_2$: Given a CNF formula $\psi(x, y)$, determine whether it is the case that for every truth assignment to x there exists a truth assignment to y such that the two assignments satisfy ψ .

We can generalize Theorem 7.2 to a broad set of schemas, by using fact-wise reductions from S^6 . This is done in the following theorem.

THEOREM 7.3. *Let $S = (\mathcal{R}, \Delta)$ be a schema such that \mathcal{R} consists of a single relation symbol R and Δ consists of two nontrivial FDs $X \rightarrow Y$ and $W \rightarrow Z$. Suppose that each of W and Z contains an attribute that is in none of the other three sets. Then g-categoricity $\langle S \rangle$ is Π_2^P -complete.*

As an example, recall that in S^6 we have $\Delta = \{\emptyset \rightarrow A, B \rightarrow C\}$. This schema is a special case of Theorem 7.3, since we can use $\emptyset \rightarrow A$ as $X \rightarrow Y$ and $B \rightarrow C$ as $W \rightarrow Z$; and indeed, each of W and Z contains an attribute (namely B and C , respectively) that is not in any of the other three sets. Additional examples of sets of FDs that satisfy the conditions of Theorem 7.3 (and hence the corresponding g-categoricity $\langle S \rangle$ is Π_2^P -complete) follow. All of these sets are over a relation symbol $R/4$. (And in each of these sets, the first FD corresponds to $X \rightarrow Y$ and the second to $W \rightarrow Z$.)

- $A \rightarrow B, C \rightarrow D$
- $A \rightarrow C, AB \rightarrow CD$
- $A \rightarrow B, ABC \rightarrow D$
- $A \rightarrow B, C \rightarrow ABD$

Unlike S^6 , to this day we do not know what is the complexity of g-categoricity $\langle S^i \rangle$ for any of the other S^i (defined in Section 5.2.2). This includes S^0 , for which all we know is membership in coNP (as stated in Theorem 4.4). However, except for this open problem, the proof technique of Theorem 5.1 is valid for g-categoricity $\langle S \rangle$. Consequently, we can show the following.

THEOREM 7.4. *The following are equivalent.*

- g-categoricity $\langle S^0 \rangle$ is coNP-hard.
- g-categoricity $\langle S \rangle$ is coNP-hard for every schema S that falls outside the polynomial-time cases of Theorem 7.1.

7.3 Transitive Priority

Let (I, \mathcal{H}, \succ) be an inconsistent prioritizing instance. We say that \succ is *transitive* if for every two facts f and g in I , if f and g are neighbors in \mathcal{H} and $f \succ^+ g$, then $f \succ g$. Transitivity is a natural assumption when \succ is interpreted as a partial order such as “is of better quality than” or “is more current than.” In this section we consider g-categoricity in the presence of this assumption. The following example shows that a g-repair is not necessarily a c-repair, even if \succ is transitive. This example provides an important context for the results that follow.

EXAMPLE 7.5. Consider again I and \succ from Example 6.4 (depicted in Figure 7). Observe that \succ is transitive. In particular, there is no priority between Ab2 and Ba2, even though $Ab2 \succ Ba2$, because Ab2 and Ba2 are not in conflict (or put differently, they are not neighbors in $\mathcal{H}_{S^1}^I$). Consider the following subinstance of I .

$$J \stackrel{\text{def}}{=} \{Aa1, Ba2, Ab3, Bb1\}$$

The reader can verify that J is a g-repair, but not a c-repair (since no execution of FindCRep can generate J). \square

Example 7.5 shows that the notion global optimality is different from completion optimality, even if the priority relation is transitive. Yet, quite remarkably, the two notions behave the same when it comes to categoricity.

THEOREM 7.6. *Let $D = (I, \mathcal{H}, \succ)$ be an inconsistent prioritizing instance such that \succ is transitive. $|\text{CRep}(D)| = 1$ if and only if $|\text{GRep}(D)| = 1$.*

PROOF. The “if” direction follows from Proposition 2.8, since every c-repair is also a g-repair. The proof of the “only if” direction is based on the special structure of the c-repair, as established in Section 6, in the case where only one c-repair exists. Specifically, suppose that there is a single c-repair J and let $J' \neq J$ be a consistent subinstance of I . We need to show that J' has a global improvement. We claim that J is a global improvement of J' . This is clearly the case if $J' \subseteq J$. So suppose that $J' \not\subseteq J$. Let f' be a fact in $J' \setminus J$. We need to show that there is a fact $f \in J \setminus J'$ such that $f \succ f'$. We complete the proof by finding such an f .

Recall from Theorem 6.5 that J is the result of executing $\text{CCategoricity}(I, \mathcal{H}, \succ)$. Consider the positive strata P_i and the negative strata N_j constructed in that execution. Since J is the union of the positive strata, we get that f' necessarily belongs to a negative stratum, say N_j . From the definition of N_j it follows that \mathcal{H} has a hyperedge e such that $f' \in e$, $(e \setminus \{f'\}) \subseteq P_1 \cup \dots \cup P_j$, and $(e \setminus \{f'\}) \succ^+ f'$. Let e be such a hyperedge. Since J' is consistent, it cannot be the case that J' contains all the facts in e . Choose a fact $f \in e$ such that $f \notin J'$. Then $f \succ^+ f'$, and since \succ is transitive (and f and f' are neighbors), we have $f \succ f'$. So $f \in J \setminus J'$ and $f \succ f'$, as required. \square

Combining Theorems 6.1 and 7.6, we get the following.

COROLLARY 7.7. *For transitive priority relations, problems g-categoricity and c-categoricity coincide, and in particular, g-categoricity is solvable in polynomial time.*

COMMENT 7.8. The reader may wonder whether Theorem 7.6 and Corollary 7.7 hold for p-categoricity as well. This is not the case. The hardness of p-categoricity(\mathbf{S}^6) (Theorem 5.7) is proved by constructing a reduction where the priority relation is transitive (and in fact, it has no chains of length larger than one). \square

In their analysis, Fagin et al. [12] have constructed various reductions for proving coNP-hardness of g-repair checking. In several of these, the priority relation is transitive. We conclude that there are schemas \mathbf{S} such that, on transitive priority relations, g-repair checking is coNP-complete whereas g-categoricity is solvable in polynomial time.

8. RELATED WORK ON DATA CLEANING

We now discuss the relationship between our work and past work on data cleaning. Specifically, we focus on relating and contrasting our complexity results with ones established in past research. To the best of our knowledge,

there has not been any work on the complexity of categoricity within the prioritized repairing of Staworko et al. [36]. Fagin et al. [13] investigated a static version of categoricity in the context of text extraction, but the settings and problems are very different, and so are the complexity results (e.g., Fagin et al. [13] establish undecidability results).

Bohannon et al. [6] have studied a repairing framework where repairing operations involve attribute updates and tuple insertions, and where the quality of a repair is determined by a cost function (aggregating the costs of individual operations). They have shown that finding an optimal repair is NP-hard, in data complexity, even when integrity constraints consist of only FDs. This result could be generalized to hardness of categoricity in their model (e.g., by a reduction from the *unique exact 3-cover* problem [35]). The source of hardness in their model is the cost minimization, and it is not clear how any of our hardness results could derive from those, as the framework of preferred repairs (adopted here) does not involve any cost-based quality; in particular, as echoed in this paper, an optimal repair can be found in polynomial time under each of the three semantics [37].

In the framework of *data currency* [15, Chapter 6] [16], relations consist of entities with attributes, where each entity may appear in different tuples, every time with possibly different (conflicting) attribute values. A partial order on each attribute is provided, where “greater than” stands for “more current.” A *completion* of an instance is obtained by completing the partial order on an attribute of every entity, and it defines a *current instance* where each attribute takes its most recent value. In addition, a completion needs to satisfy given (denial) constraints, which may introduce interdependencies among completions of different attributes. Fan et al. [16] have studied the problem of determining whether such a specification induces a single current instance (i.e., the corresponding version of categoricity), and showed that this problem is coNP-complete under data complexity. It is again not clear how to simulate their hardness in our p-categoricity and g-categoricity, since their hardness is due to the constraints on completions, and these constraints do not have correspondents in our case (beyond the partial orders). A similar argument relates our lower bounds to those in the framework of conflict resolution by Fan Geerts [15, Chapter 7.3], where the focus is on establishing a unique tuple from a collection of conflicting tuples.

Fan et al. [16] show that in the absence of constraints, their categoricity problem can be solved in polynomial time (even in the presence of “copy functions”). This tractability result can be used for establishing the tractability side of Theorem 5.1 in the special case where the single FD is a key constraint. In the general case of a single FD, we need to argue about relationships among sets, and in particular, the differences among the three x-categoricity problems matter.

Cao et al. [8] have studied the problem of entity record cleaning, where again the attributes of an entity are represented as a relation (with missing values), and a partial order is defined on each attribute. The goal is to increase the accuracy of values from the partial orders and an external source of reliable (“master”) data. The specification now gives up-

date steps that have the form of logical rules that specify when one value should replace a null, when new preferences are to be derived, and when data should be copied from the master data. Hence, cleaning is established by *chasing* these rules. They study a problem related to categoricity, namely the *Church-Rosser* property: is it the case that every application of the chase (in any rule-selection and grounding order) results in the same instance? They show that this property is testable in polynomial time by giving an algorithm that tests whether some invalid step in the end of the execution has been valid sometime during the execution. We do not see any clear way of deriving any of our upper bounds from this result, due to the difference in the update model (updating nulls and preferences vs. tuple deletion), and the optimality model (chase termination vs. x-repair).

The works on *certain fixes* [17, 18] [15, Chapters 7.1–7.2] consider models that are substantially different from the one adopted here, where repairs are obtained by chasing update rules (rather than tuple deletion), and uniqueness applies to chase outcomes (rather than maximal subinstances w.r.t. preference lifting). The problems relevant to our categoricity are the *consistency problem* [18] (w.r.t. guarantees on the consistency of some attributes following certain patterns), and the *determinism problem* [18]. They are shown to be intractable (coNP-complete [17] and PSPACE-complete [18]) under combined complexity (while we focus here on data complexity).

Finally, we remark that there have several dichotomy results on the complexity of problems associated with inconsistent data [12, 28, 31], but to the best of our knowledge this paper is the first to establish a dichotomy result for any variant of repair uniqueness identification.

9. CONCLUDING REMARKS

We investigated the complexity of the categoricity problem, which is that of determining whether the provided priority relation suffices to clean the database unambiguously in the framework of preferred repairs. Following the three semantics of optimal repairs, we investigated the three variants of this problem: p-categoricity, g-categoricity and c-categoricity. We established a dichotomy in the data complexity of p-categoricity for the case where constraints are FDs, partitioning the cases into polynomial time and coNP-completeness. We further showed that the tractable side of p-categoricity extends to g-categoricity, but the latter can reach Π_2^P -completeness already for two FDs. Finally, we showed that c-categoricity is solvable in polynomial time in the general case where integrity constraints are given as a conflict hypergraph. We complete this paper by discussing directions for future research.

In this work we did not address any qualitative discrimination among the three notions of x-repairs. Rather, we continue the line of work [13, 36] that explores the impact of the choice on the entailed computational complexity. It has been established that, as far as repair checking is concerned, the Pareto and the completion semantics behave much better than the global one, since g-repair checking is tractable

only in a very restricted class of schemas [13]. In this work we have shown that from the viewpoint of categoricity, the Pareto semantics departs from the completion one by being likewise intractable (while the global semantics hits an even higher complexity class), hence the completion semantics outstands so far as the most efficient option to adopt.

It would be interesting to further understand the complexity of g-categoricity, towards a dichotomy (at least for FDs). We have left open the question of whether there exists a schema with a single relation and a set of FDs, *not* equivalent to a single FD, such that g-categoricity is solvable in polynomial time. Beyond that, for both p-categoricity and g-categoricity it is important to detect islands of tractability based on properties of the data and/or the priority relation (as schema constraints do not get us far in terms of efficient algorithms, at least by our dichotomy for p-repairs), beyond transitivity in the case of g-categoricity (Corollary 7.7).

Another interesting direction would be the generalization of categoricity to the problems of *counting* and *enumerating* the preferred repairs. For classical repairs (without a priority relation), Maslowski and Wijsen [31, 32] established dichotomies (FP vs. #P-completeness) in the complexity of counting in the case where constraints are primary keys. For the general case of denial constraints, counting the classical repairs reduces to the enumeration of independent sets of a hypergraph with a bounded edge size, a problem shown by Boros et al. [7] to be solvable in incremental polynomial time (and in particular polynomial input-output complexity). For a general given conflict hypergraph, repair enumeration is the well known problem of enumerating the *minimal hypergraph transversals* (also known as the hypergraph *duality* problem); whether this problem is solvable in polynomial total time is a long standing open problem [22].

In this work we focused on cleaning within the framework of preferred repairs, where integrity constraints are anti-monotonic and cleaning operations are tuple deletions (i.e., *subset repairs*). However, the problem of categoricity arises in every cleaning framework that is based on defining a set of repairs with a preference between repairs, including different types of integrity constraints, different cleaning operations (e.g., tuple addition and cell update [38]), and different priority specifications among repairs. This includes preferences by means of general scoring functions [24, 33], aggregation of scores on the individual cleaning operations [6, 11, 18, 18, 27], priorities among resolution policies [30] and preferences based on soft rules [23, 34]. This also includes the LLUNATIC system [20, 21] where priorities are defined by lifting partial orders among “cell groups,” representing either semantic preferences (e.g., timestamps) or level of completeness (e.g., null vs. non-null). A valuable future direction would be to investigate the complexity of categoricity in the above frameworks, and in particular, to see whether ideas or proof techniques from this work can be used to analyze their categoricity.

Motivated by the tractability of c-categoricity, we plan to pursue an implementation of an interactive and declarative system for database cleaning, where rules are of two kinds: integrity constraints and priority specifications (e.g., based

on the semantics of *priority generating dependencies* of Fagin et al. [13]). To make such a system applicable to a wide range of practical use cases, we will need to extend beyond subset repairs, and consequently, investigate the fundamental direction of extending the framework of preferred repairs towards such repairs.

10. REFERENCES

- [1] F. N. Afrati and P. G. Kolaitis. Repair checking in inconsistent databases: algorithms and complexity. In *ICDT*, pages 31–41. ACM, 2009.
- [2] D. E. Appelt and B. Onyshkevych. The common pattern specification language. In *TIPSTER Text Program: Phase III*, pages 23–30. Association for Computational Linguistics, 1998.
- [3] M. Arenas, L. E. Bertossi, and J. Chomicki. Consistent query answers in inconsistent databases. In *PODS*, pages 68–79. ACM, 1999.
- [4] L. E. Bertossi. *Database Repairing and Consistent Query Answering*. Synthesis Lectures on Data Management. Morgan & Claypool Publishers, 2011.
- [5] P. Bohannon, W. Fan, F. Geerts, X. Jia, and A. Kementsietsidis. Conditional functional dependencies for data cleaning. In *ICDE*, pages 746–755. IEEE, 2007.
- [6] P. Bohannon, M. Flaster, W. Fan, and R. Rastogi. A cost-based model and effective heuristic for repairing constraints by value modification. In *SIGMOD*, pages 143–154. ACM, 2005.
- [7] E. Boros, K. M. Elbassioni, V. Gurvich, and L. Khachiyan. An efficient incremental algorithm for generating all maximal independent sets in hypergraphs of bounded dimension. *Parallel Processing Letters*, 10(4):253–266, 2000.
- [8] Y. Cao, W. Fan, and W. Yu. Determining the relative accuracy of attributes. In *SIGMOD*, pages 565–576. ACM, 2013.
- [9] L. Chiticariu, R. Krishnamurthy, Y. Li, S. Raghavan, F. Reiss, and S. Vaithyanathan. SystemT: An algebraic approach to declarative information extraction. In *ACL*, pages 128–137, 2010.
- [10] J. Chomicki and J. Marcinkowski. Minimal-change integrity maintenance using tuple deletions. *Inf. Comput.*, 197(1-2):90–121, 2005.
- [11] M. Dallachiesa, A. Ebaid, A. Eldawy, A. K. Elmagarmid, I. F. Ilyas, M. Ouzzani, and N. Tang. NADEEF: a commodity data cleaning system. In *SIGMOD*, pages 541–552. ACM, 2013.
- [12] R. Fagin, B. Kimelfeld, and P. G. Kolaitis. Dichotomies in the complexity of preferred repairs. In *PODS*, pages 3–15. ACM, 2015.
- [13] R. Fagin, B. Kimelfeld, F. Reiss, and S. Vansummeren. Cleaning inconsistencies in information extraction via prioritized repairs. In *PODS*, pages 164–175. ACM, 2014.
- [14] R. Fagin, B. Kimelfeld, F. Reiss, and S. Vansummeren. Document spanners: A formal approach to information extraction. *J. ACM*, 62(2):12, 2015.
- [15] W. Fan and F. Geerts. *Foundations of Data Quality Management*. Synthesis Lectures on Data Management. Morgan & Claypool Publishers, 2012.
- [16] W. Fan, F. Geerts, and J. Wijsen. Determining the currency of data. *ACM Trans. Database Syst.*,

- 37(4):25, 2012.
- [17] W. Fan, J. Li, S. Ma, N. Tang, and W. Yu. Towards certain fixes with editing rules and master data. *VLDB J.*, 21(2):213–238, 2012.
 - [18] W. Fan, S. Ma, N. Tang, and W. Yu. Interaction between record matching and data repairing. *J. Data and Information Quality*, 4(4):16:1–16:38, 2014.
 - [19] T. Gaasterland, P. Godfrey, and J. Minker. An overview of cooperative answering. *J. Intell. Inf. Syst.*, 1(2):123–157, 1992.
 - [20] F. Geerts, G. Mecca, P. Papotti, and D. Santoro. The LLUNATIC data-cleaning framework. *PVLDB*, 6(9):625–636, 2013.
 - [21] F. Geerts, G. Mecca, P. Papotti, and D. Santoro. Mapping and cleaning. In *ICDE*, pages 232–243. IEEE, 2014.
 - [22] G. Gottlob and E. Malizia. Achieving new upper bounds for the hypergraph duality problem through logic. In *CSL-LICS*, pages 43:1–43:10. ACM, 2014.
 - [23] S. Greco, C. Sirangelo, I. Trubitsyna, and E. Zumpano. Feasibility conditions and preference criteria in querying and repairing inconsistent databases. In *DEXA*, volume 3180 of *LNCS*, pages 44–55. Springer, 2004.
 - [24] S. Greco, C. Sirangelo, I. Trubitsyna, and E. Zumpano. Preferred repairs for inconsistent databases. In *Encyclopedia of Database Technologies and Applications*, pages 480–485. Idea Group, 2005.
 - [25] B. Kimelfeld. A dichotomy in the complexity of deletion propagation with functional dependencies. In *PODS*, pages 191–202, 2012.
 - [26] B. Kimelfeld, J. Vondrák, and R. Williams. Maximizing conjunctive views in deletion propagation. *ACM Trans. Database Syst.*, 37(4):24, 2012.
 - [27] S. Kolahi and L. V. S. Lakshmanan. On approximating optimum repairs for functional dependency violations. In *ICDT*, volume 361 of *ACM International Conference Proceeding Series*, pages 53–62. ACM, 2009.
 - [28] P. Koutris and D. Suciu. A dichotomy on the complexity of consistent query answering for atoms with simple keys. In *ICDT*, pages 165–176. OpenProceedings.org, 2014.
 - [29] P. Koutris and J. Wijsen. The data complexity of consistent query answering for self-join-free conjunctive queries under primary key constraints. In *PODS*, pages 17–29. ACM, 2015.
 - [30] M. V. Martinez, F. Parisi, A. Pugliese, G. I. Simari, and V. S. Subrahmanian. Inconsistency management policies. In *KR*, pages 367–377. AAAI Press, 2008.
 - [31] D. Maslowski and J. Wijsen. A dichotomy in the complexity of counting database repairs. *J. Comput. Syst. Sci.*, 79(6):958–983, 2013.
 - [32] D. Maslowski and J. Wijsen. Counting database repairs that satisfy conjunctive queries with self-joins. In *ICDT*, pages 155–164. OpenProceedings.org, 2014.
 - [33] A. Motro, P. Anokhin, and A. C. Acar. Utility-based resolution of data inconsistencies. In *IQIS*, pages 35–43. ACM, 2004.
 - [34] D. V. Nieuwenborgh and D. Vermeir. Preferred answer sets for ordered logic programs. *TPLP*, 6(1-2):107–167, 2006.
 - [35] O. Schulte, W. Luo, and R. Greiner. Mind change optimal learning of bayes net structure from dependency and independency data. *Inf. Comput.*, 208(1):63–82, 2010.
 - [36] S. Staworko, J. Chomicki, and J. Marcinkowski. Preference-driven querying of inconsistent relational databases. In *EDBT Workshops*, volume 4254 of *LNCS*, pages 318–335. Springer, 2006.
 - [37] S. Staworko, J. Chomicki, and J. Marcinkowski. Prioritized repairing and consistent query answering in relational databases. *Ann. Math. Artif. Intell.*, 64(2-3):209–246, 2012.
 - [38] J. Wijsen. Database repairing using updates. *ACM Trans. Database Syst.*, 30(3):722–768, 2005.

APPENDIX

A. PROOFS FOR SECTION 5

In this section we provide proofs for Section 5.

A.1 Proof of Lemma 5.3

In the following section, we say that a block I_a (respectively, $I_{a,b}$) is the block (respectively, subblock) of a fact f if $f \in I_a$ (respectively $f \in I_{a,b}$). Note that each fact has a unique block and subblock.

LEMMA 5.3. *Let J be a subinstance of I . Then J is a p-repair if and only if J is a union of p-repairs over all the blocks I_a of I . Moreover, each p-repair of a block I_a is a subblock $I_{a,b}$.*

PROOF. Recall that Δ is the set $\{A \rightarrow B\}$. We start by proving the second part of the lemma. That is, we show that each p-repair of a block I_a is a subblock $I_{a,b}$. Let K be a p-repair of I_a . Then K is contained in a single subblock of I_a , since K is consistent. Moreover, K contains all the facts in $I_{a,b}$, or otherwise K has a Pareto improvement.

Next, we prove the first part of the lemma.

The “if” direction. Let J be a p-repair of I . We need to show that J is a union of p-repairs over all the blocks I_a of I . Observe that J is consistent, and so, for each block I_a it contains facts from at most one subblock $I_{a,b}$. Moreover, since J is maximal, it contains at least one representative from each block I_a , and furthermore, it contains the entire subblock of each such a representative. We conclude that J is the union of subblocks of I . It is left to show that if a subblock $I_{a,b}$ is contained in J , then $I_{a,b}$ is a p-repair of I_a . Let $I_{a,b}$ be a subblock contained in J and assume, by way of contradiction, that K is a Pareto-improvement of $I_{a,b}$ in I_a . Let J' be the instance that is obtained from J by replacing $I_{a,b}$ with K . Observe that J' is consistent, since no facts in J other than those in $I_{a,b}$ conflict with facts from K . Then clearly, J' is a Pareto improvement of J , which contradicts the fact that J is a p-repair.

The “only if” direction. Let J be a union of p-repairs over all the blocks I_a of I . We need to show that J is a p-repair. By the second part of the lemma, J is a union of subblocks. Since each subblock is consistent and facts from different blocks are consistent, we get that J is consistent. It is left to show that J does not have a Pareto improvement. Assume, by way of contradiction, that J has a Pareto improvement K . By the definition of a Pareto improvement, K contains a fact f such that $f \succ g$ for all $g \in J \setminus K$. Let f be such a fact. Let I_a be the subblock of f . Then, by our assumption the subinstance J contains a p-repair of I_a , and from the second part of the lemma this p-repair is a subblock of I_a , say $I_{a,b}$. But then, f is not in $I_{a,b}$ (since $f \notin J$), and therefore, K does not contain any fact from $I_{a,b}$ (since K is consistent). We conclude that $f \succ g$ for all $g \in I_{a,b}$, and hence, $I_{a,b}$ has a Pareto improvement (namely $\{f\}$), in contradiction to the fact that $I_{a,b}$ is a p-repair of I_a . \square

A.2 Proof of Theorem 5.7

We now prove Theorem 5.7

THEOREM 5.7. *The problems $p\text{-categoricity}\langle S^0 \rangle$ and $p\text{-categoricity}\langle S^6 \rangle$ are both coNP-hard.*

We give a separate proof for each of the two schemas.

A.2.1 Hardness of $p\text{-categoricity}\langle S^0 \rangle$

We construct a reduction from the Exact-Cover problem (XC) to the complement of $p\text{-categoricity}\langle S^0 \rangle$. The input to XC is a set \mathcal{U} of elements and a collection \mathcal{X} of subsets of \mathcal{U} , such that their union is \mathcal{U} . The goal is to identify whether there is an exact cover of \mathcal{U} by \mathcal{X} . An *exact cover* of \mathcal{U} by \mathcal{X} is a collection of pairwise disjoint sets from \mathcal{X} whose union is \mathcal{U} .

Construction. Given an input $(\mathcal{X}, \mathcal{U})$ to XC, we construct input (I, \succ) for $p\text{-categoricity}\langle S^0 \rangle$. For each $u \in \mathcal{U}$, $X \in \mathcal{X}$ and $x \in X$, the instance I consists of the following facts:

- (i) $R^0(u, u)$
- (ii) $R^0(u, f_u)$
- (iii) $R^0(X_x, f_u)$
- (iv) $R^0(X_x, X_x)$
- (v) $R^0(X_x, x)$
- (vi) $R^0(X_{x_{i+1}}, x_i)$ for each $i = 0, \dots, n-1$, where $X = \{x_0, \dots, x_{n-1}\}$ and plus is interpreted modulo n (e.g., $(n-1)+1 = 0$)

In the sequel, we relate to these facts by *types* according to their roman number. For example, facts of the form $R^0(X_x, f_u)$, where $X \in \mathcal{X}$, $x \in X$ and $u \in \mathcal{U}$, will be referred to as facts of type (iii).

For all $u \in \mathcal{U}$, $X \in \mathcal{X}$ and $x \in X$, the priority relation \succ is defined as follows:




Figure 8: Illustration of the reduction from XC to p-categoricity(S^0)

- $R^0(u, u) \succ R^0(u, f_u)$,
- $R^0(X_x, X_x) \succ R^0(X_x, f_u)$,
- $R^0(u, u) \succ R^0(X_x, u)$,
- $R^0(X_x, f_u) \succ R^0(X_x, x)$, and
- $R^0(X_{x_{i+1}}, x_i) \succ R^0(X_{x_i}, x_i)$, for $i = 0, \dots, n-1$ where $X = \{x_0, \dots, x_{n-1}\}$ and plus is interpreted modulo n . Note that each $X \in \mathcal{X}$ has a corresponding n .

Our construction is partly illustrated in Figure 8 for the following input to the XC problem: $\mathcal{U} = \{1, 2, 3, 4\}$, $\mathcal{X} = \{A, B, C\}$ where $A = \{1, 2\}$, $B = \{3, 4\}$ and $C = \{2, 3\}$. Note that we denote the fact $R^0(a, b)$ by (a, b) . In this case, there is an exact cover of \mathcal{U} by \mathcal{X} that consists of the sets A and B . The gray facts represent a p-repair.

Proof of Hardness. We start by finding a c-repair.

LEMMA A.1. *There is a c-repair that consists of the following facts for all $u \in \mathcal{U}$, $X \in \mathcal{X}$ and $x \in X$.*

- $R^0(u, u)$
- $R^0(X_x, X_x)$

PROOF. It is straightforward to show that there is a run of the algorithm $\text{FindCRep}(I, \mathcal{H}, \succ)$ that returns exactly this c-repair. \square

In the remainder of the proof, we relate to the c-repair from Lemma A.1 by J_0 . Note that every c-repair is also a p-repair and therefore J_0 is a p-repair of I . To complete the proof we show that there is a solution to XC if and only if I has a p-repair different from J_0 .

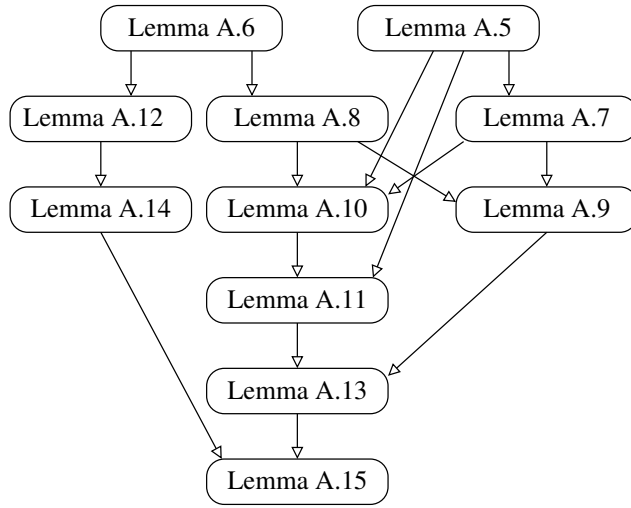


Figure 9: Dependencies between the lemmas in the proof of the “only if” direction of p-categoricity $\langle S^0 \rangle$

The “if” direction. We show that if there is a solution to XC then I has a p-repair different from J_0 .

We construct a p-repair of I , namely K , different from J_0 based on a solution to XC. Let the collection of sets $X^1, \dots, X^l \in \mathcal{X}$ be a solution to XC. Let K consist of the following facts for all $X \in \mathcal{X}$, $x \in X$ and $u \in \mathcal{U}$.

- (1) $R^0(X_x, x)$ if $X \in \{X^1, \dots, X^l\}$
- (2) $R^0(X_x, X_x)$ if $X \notin \{X^1, \dots, X^l\}$
- (3) $R^0(u, f_u)$

Note that since K is different from J_0 (see Lemma A.1), it is left to show that K is a p-repair of I . To do so, we show that K is consistent and that it does not have a Pareto improvement.

LEMMA A.2. K is a consistent subinstance of I .

The proof of this lemma is straightforward based on the above construction.

LEMMA A.3. K does not have a Pareto improvement.

PROOF. It suffices to show that for all f in $I \setminus K$, there exists f' in K such that $\{f, f'\}$ is inconsistent (w.r.t. Δ^0) and $f \not\preceq f'$. For each $f \in I \setminus K$ we choose $f' \in K$ such that the conditions hold. We divide to different cases according to the type of f .

- **f is of type (i):** That is, there exists an element u in \mathcal{U} such that $f = R^0(u, u)$. Since the collection $\{X^1, \dots, X^l\}$ is a cover of \mathcal{U} , there exists X in $\{X^1, \dots, X^l\}$ such that $u \in X$. Hence, $R^0(X_u, u) \in K$ and we choose $f' = R^0(X_u, u)$.
- **f is of type (ii):** This is impossible, since K contains all the facts of this type.
- **f is of type (iii):** That is, there exists $X \in \mathcal{X}$, $x \in X$ and $u \in \mathcal{U}$ such that $f = R^0(X_x, f_u)$. We choose $f' = R^0(u, f_u)$.
- **f is of type (iv):** That is, there exists a set $X \in \mathcal{X}$ and $x \in X$ such that $f = R^0(X_x, X_x)$. Since $f \notin K$, it holds that $X \notin \{X^1, \dots, X^l\}$. Hence $R^0(X_x, x) \in K$ and we choose $f' = R^0(X_x, x)$.
- **f is of type (v):** That is, there exists a set $X \in \mathcal{X}$ and $x \in X$ such that $f = R^0(X_x, x)$. Since $f \notin K$, it holds that $X \notin \{X^1, \dots, X^l\}$. Hence $R^0(X_x, X_x) \in K$ and we choose $f' = R^0(X_x, X_x)$.
- **f is of type (vi):** That is, there exists a set $X \in \mathcal{X}$ where $X = \{x_0, \dots, x_{n-1}\}$ and $0 \leq i \leq n-1$ such that $f = R^0(X_{x_{i+1}}, x_i)$. Since the collection $\{X^1, \dots, X^l\}$ is a cover of \mathcal{U} , there exists some Y in $\{X^1, \dots, X^l\}$ such that $x_i \in Y$. Hence, we get that $R^0(Y_{x_i}, x_i) \in K$ and we choose $f' = R^0(Y_{x_i}, x_i)$.

It holds that in all of the above cases $\{f, f'\}$ is inconsistent and $f \not\preceq f'$. \square

Lemmas A.2 and A.3 imply the following.

LEMMA A.4. K is a p-repair of I .

Given a solution to XC, we constructed a p-repair K different from the above c-repair. This completes the “if” direction.

The “only if” direction. We show that if I has a p-repair different from J_0 (from Lemma A.1) then there is a solution to XC. The proof of this direction consists of several lemmas, and the dependencies between them are described in Figure 9. For example, the proof of Lemma A.10 is based on Lemmas A.5, A.7 and A.8.

Let K be a p-repair different from J_0 . In the next two lemmas we show that facts of types (iii) and (vi) are not in K .

LEMMA A.5. *For all $X \in \mathcal{X}$, $x \in X$ and $u \in \mathcal{U}$, it holds that $R^0(X_x, f_u) \notin K$.*

PROOF. Let $X \in \mathcal{X}$, $x \in X$ and $u \in \mathcal{U}$. Assume, by way of contradiction, that $R^0(X_x, f_u) \in K$. Since K is consistent and $R^0(X_x, X_x)$ is inconsistent with $R^0(X_x, f_u)$, we obtain that $R^0(X_x, X_x) \notin K$. Since $R^0(X_x, X_x) \succ R^0(X_x, f_u)$, it holds that K must contain a fact f that is inconsistent with $R^0(X_x, X_x)$ (but is consistent with $R^0(X_x, f_u)$). Therefore, f must agree with $R^0(X_x, X_x)$ on B . That leads to a contradiction since there is no such fact. \square

LEMMA A.6. *For all $X \in \mathcal{X}$, $x \in X$ and $u \in \mathcal{U}$ where $u \neq x$, it holds that $R^0(X_x, u) \notin K$.*

PROOF. Let $X \in \mathcal{X}$, $x \in X$ and $u \in \mathcal{U}$ where $u \neq x$. Assume that $R^0(X_x, u) \in K$. Since for all $X \in \mathcal{X}$, $x \in X$ and $u \in \mathcal{U}$, we have that $R^0(u, u) \succ R^0(X_x, u)$, the p-repair K must contain a fact f that is inconsistent with $R^0(u, u)$ (but is consistent with $R^0(X_x, u)$). Since K is consistent, f must agree with $R^0(u, u)$ on A . Thus, f must be the fact $R^0(u, f_u)$. Replacing both facts $R^0(X_x, u)$ and $R^0(u, f_u)$ with $R^0(u, u)$ results in a Pareto improvement of K which leads to a contradiction since K is a p-repair. \square

We establish a connection between facts of type (i) and (ii).

LEMMA A.7. *Let $u \in \mathcal{U}$. If the fact $R^0(u, u) \notin K$ then $R^0(u, f_u) \in K$.*

PROOF. Assume $R^0(u, u) \notin K$. Assume, by way of contradiction, that $R^0(u, f_u) \notin K$. Since K is maximal, it must contain a fact f that is inconsistent with $R^0(u, f_u)$. If f agrees with $R^0(u, f_u)$ on A then it can only be of type (i). That is, the only possibility is that $f = R^0(u, u)$ which leads to a contradiction. If f agrees with $R^0(u, f_u)$ on B , it can only be of type (iii) which leads to a contradiction since by Lemma A.5, such a fact cannot be in a p-repair. \square

Moreover, we establish a connection between facts of type (ii) and (v).

LEMMA A.8. *Let $u \in \mathcal{U}$. If $R^0(u, f_u) \in K$ then there exists $X \in \mathcal{X}$ such that $u \in X$ and $R^0(X_u, u) \in K$.*

PROOF. Assume $R^0(u, f_u) \in K$. Since K is consistent, $R^0(u, u) \notin K$. It follows from $R^0(u, u) \succ R^0(u, f_u)$ that K must contain a fact f that is inconsistent with $R^0(u, u)$. Since K is consistent, f is consistent with $R^0(u, f_u)$. Therefore, f must agree with $R^0(u, u)$ on B . The possible types for f are (vi) and (v). By Lemma A.6, f is not of type (vi). Thus, f must be of type (v). Therefore, $R^0(X_u, u) \in K$ for $X \in \mathcal{X}$ such that $u \in X$ (there exists such X since the union of sets of \mathcal{X} is \mathcal{U}). \square

We conclude the following connection between facts of types (i) and (v).

LEMMA A.9. *Let $u \in \mathcal{U}$. If $R^0(X_u, u) \notin K$ for all $X \in \mathcal{X}$ such that $u \in X$, then $R^0(u, u) \in K$.*

PROOF. Assume $R^0(X_u, u) \notin K$ for all $X \in \mathcal{X}$. Assume, by way of contradiction, that $R^0(u, u) \notin K$. By Lemma A.7, $R^0(u, f_u) \in K$. By Lemma A.8, there exists $X \in \mathcal{X}$ such that $u \in X$ and $R^0(X_u, u) \in K$. This is a contradiction. \square

We state that either $R^0(u, u) \in K$ for all $u \in \mathcal{U}$ or none of the facts $R^0(u, u)$, where $u \in \mathcal{U}$, is in K .

LEMMA A.10. *Let $u_0 \in \mathcal{U}$. If $R^0(u_0, u_0) \in K$ then $R^0(u, u) \in K$ for all $u \in \mathcal{U}$.*

PROOF. Let $u_0 \in \mathcal{U}$ and assume $R^0(u_0, u_0) \in K$. Assume, by way of contradiction, there exists $u \in \mathcal{U}$ such that $u \neq u_0$ and $R^0(u, u) \notin K$. By Lemma A.7, $R^0(u, f_u) \in K$. Thus, by Lemma A.8 we have that $R^0(U_u, u) \in K$ for some $U \in \mathcal{X}$ such that $u \in U$. Note that since $R^0(U_u, f_{u_0}) \succ R^0(U_u, u)$, there must be a fact in K that is inconsistent with $R^0(U_u, f_{u_0})$ and is consistent with $R^0(U_u, u)$. The only such a fact is $R^0(u_0, f_{u_0})$ (i.e., $R^0(u_0, f_{u_0}) \in K$). This is a contradiction since $R^0(u_0, u_0) \in K$ and K is consistent. \square

We prove that all facts of type (i) are not in K .

LEMMA A.11. *For all $u \in \mathcal{U}$, $R^0(u, u) \notin K$.*

PROOF. Let $u \in \mathcal{U}$ and assume, by way of contradiction, that $R^0(u, u) \in K$. By Lemma A.10, for all $u \in \mathcal{U}$ $R^0(u, u) \in K$. Since K is consistent, it does not contain facts of the form $R^0(u, f_u)$ for all $u \in \mathcal{U}$ (since the fact $R^0(u, f_u)$ is inconsistent with $R^0(u, u)$). Moreover, K does not contain facts of the form $R^0(X_u, u)$ where $X \in \mathcal{X}$ and $u \in X$ (for a similar reason). By Lemma A.5 (respectively, A.6), K does not contain facts of the form $R^0(X_x, f_u)$ (respectively, $R^0(X_x, u)$) where $X \in \mathcal{X}$ and $u \in X$. It holds that K is maximal and thus it contains all of the facts of type (iv). Thus, we conclude that K is exactly J_0 which leads to a contradiction. \square

We show that if a fact of type (iv) is in K , then all of the facts of type (iv) are in K .

LEMMA A.12. *Let $X \in \mathcal{X}$. If $R^0(X_{x'}, x') \in K$ for some $x' \in X$ then for all $x \in X$ we have $R^0(X_x, x) \in K$.*

PROOF. Let us denote $X = \{x_0, \dots, x_{n-1}\}$ and assume without loss of generality that $x_0 = x'$. We prove by induction on i that $R^0(X_{x_i}, x_i) \in K$ for all $i = 0, \dots, n-1$.

Basis. Trivial.

Induction Step. Assume $R^0(X_{x_i}, x_i) \in K$. Since $R^0(X_{x_{i+1}}, x_i) \succ R^0(X_{x_i}, x_i)$, there must be a fact $f \in K$ that is inconsistent with $R^0(X_{x_{i+1}}, x_i)$. Since K is consistent, f must agree with $R^0(X_{x_{i+1}}, x_i)$ on A . By Lemma A.5, f is not of the form $R^0(X_{x_{i+1}}, f_u)$ for $u \in \mathcal{U}$ and by Lemma A.6 f is also not of the form $R^0(X_{x_{i+1}}, u)$ for $u \in \mathcal{U}$. Therefore, f must be $R^0(X_{x_{i+1}}, x_{i+1})$. \square

Next, we show that K encodes a solution for XC. Specifically, we contend that there is an exact cover of \mathcal{U} by \mathcal{X} , namely \mathcal{C} , that is defined as follows: $C \in \mathcal{C}$ if and only if $R^0(C_c, c) \in K$ for some $c \in C$.

LEMMA A.13. *For all $u \in \mathcal{U}$ there exists $C \in \mathcal{C}$ such that $u \in C$.*

PROOF. Let $u \in \mathcal{U}$ and assume, by way of contradiction, that for all $C \in \mathcal{C}$, it holds that $u \notin C$. By our assumption, $X \notin \mathcal{C}$ for all $X \in \mathcal{X}$ such that $u \in X$. Note that since the union of the sets in \mathcal{X} is \mathcal{U} , there exists a set $X \in \mathcal{X}$ such that $u \in X$. Thus, the definition of the set \mathcal{C} implies that for all $X \in \mathcal{X}$ such that $u \in X$, we have that $R^0(X_u, u) \notin K$. By Lemma A.9, $R^0(u, u) \in K$. By Lemma A.11, this is a contradiction. \square

LEMMA A.14. *For all $C, C' \in \mathcal{C}$ where $C \neq C'$, it holds that $C \cap C' = \emptyset$.*

PROOF. Let $C, C' \in \mathcal{C}$ where $C \neq C'$. Assume, by way of contradiction, that there exists $u \in C \cap C'$. By \mathcal{C} 's definition, $R^0(C_c, c) \in K$ for some $c \in C$. Lemma A.12 implies that for all $c \in C$, it holds that $R^0(C_c, c) \in K$. Similarly, for all $c' \in C'$, we have that $R^0(C'_{c'}, c') \in K$. Since $u \in C \cap C'$, both facts $R^0(C_u, u)$ and $R^0(C'_u, u)$ are in K . This is a contradiction to K 's consistency. \square

Finally, we conclude the following.

LEMMA A.15. *\mathcal{C} is an exact cover of \mathcal{U}*

PROOF. Follows from Lemmas A.13 and A.14. \square

Given K , a p-repair different from J_0 , we showed that there exists a solution to XC, (i.e., an exact cover of \mathcal{U}). This completes the “only if” direction.

A.2.2 Hardness of p -categoricity $\langle \mathbf{S}^6 \rangle$

We construct a reduction from CNF satisfiability to p -categoricity $\langle \mathbf{S}^6 \rangle$. The input to CNF is a formula ψ with the free variables x_1, \dots, x_n , such that ψ has the form $c_0 \wedge \dots \wedge c_m$ where each c_j is a clause. Each clause is a conjunction of variables from the set $\{x_i, \neg x_i : i = 1, \dots, n\}$. The goal is to determine whether there is a true assignment $\tau : \{x_1, \dots, x_n\} \rightarrow \{0, 1\}$ that satisfies ψ . Given such an input, we will construct the input (I, \succ) for p -categoricity $\langle \mathbf{S}^6 \rangle$. For each $i = 1, \dots, n$ and $j = 0, \dots, m$, I contains the following facts:

- $R^6(\odot, x_i, 0)$
- $R^6(\odot, x_i, 1)$
- $R^6(\otimes, c_j, c_j)$

The priority relation \succ is defined as follows:

- $R^6(\otimes, c_j, c_j) \succ R^6(\odot, x_i, 0)$ if x_i appears in clause c_j ,
- $R^6(\otimes, c_j, c_j) \succ R^6(\odot, x_i, 1)$ if $\neg x_i$ appears in clause c_j , and
- $R^6(\otimes, c_j, c_j) \succ R^6(\odot, x_i, b)$, where $b \in \{0, 1\}$, if neither x_i nor $\neg x_i$ appear in clause c_j .

Our construction is illustrated in Figure 10 for the CNF formula: $\psi = (x_1 \vee x_2 \vee \neg x_4) \wedge (\neg x_2 \vee x_3 \vee \neg x_4) \wedge (\neg x_1 \vee \neg x_3 \vee x_4)$. Observe that the subinstance J that consists of the facts $R^6(\otimes, c_j, c_j)$ for all j , is the only c-repair. To complete the proof, we will show that ψ is satisfiable if and only if I has a p-repair different from J .

The “if” direction. Assume ψ is satisfiable. That is, there exists an assignment $\tau : \{x_1, \dots, x_n\} \rightarrow \{0, 1\}$ that satisfies ψ . We claim that the subinstance K that consists of the facts $R^6(\odot, x_i, \tau(x_i))$ for all i is a p-repair (that is different from J). K is consistent since τ is an assignment (i.e., each x_i has exactly one value and the constraint $B \rightarrow C \in \Delta^6$ is satisfied). For




Figure 10: Illustration of the reduction from CNF to p-categoricity $\langle S^6 \rangle$

the same reason, K is maximal (facts of the form $R^6(\otimes, c_j, c_j)$ cannot be added to K because of the constraint $\emptyset \rightarrow A$). It is left to show that K does not have a Pareto improvement. Assume, by way of contradiction, that it does. That is, there exists a fact $f \in I \setminus K$ such that $f \succ g$ for every $g \in K$. Note that follows from \succ 's definition, f must be of the form $R^6(\otimes, c_j, c_j)$. Nevertheless, this implies the clause c_j is not satisfied by τ which leads us to the conclusion that ψ is not satisfied by τ .

The “only if” direction. Assume there is a p-repair K different from J . Since K is different from J , it must contain a fact $R^6(\odot, x_i, b_i)$ for some i . Since K is consistent and the FD $\emptyset \rightarrow A$ is in Δ^6 , it holds that for all facts f in K , we have that $f[A] = \odot$. K is maximal and consistent and thus induces a true assignment τ , defined by $\tau(x_i) \stackrel{\text{def}}{=} b_i$. Assume, by way of contradiction, that τ does not satisfy ψ . That is, there exists a clause c_j that is not satisfied. By \succ 's definition, $R^6(\otimes, c_j, c_j) \succ R^6(\odot, x_i, b_i)$ for each x_i that appears in c_j (with or without negation). Moreover, $R^6(\otimes, c_j, c_j) \succ R^6(\odot, x_l, b_l)$ for each x_l that does not appear in c_j (with or without negation). Therefore $R^6(\otimes, c_j, c_j) \succ f$ for every $f \in K$, which implies that K has a Pareto improvement that contains the fact $R^6(\otimes, c_j, c_j)$. Hence, we get a contradiction.

A.3 Proof of Lemma 5.9

LEMMA 5.9. Let $S = (\mathcal{R}, \Delta)$ be a schema such that \mathcal{R} consists of a single relation symbol. Suppose that Δ is equivalent to a pair of keys, and Δ is not equivalent to any single FD. Then there is a fact-wise reduction from S^0 to S .

PROOF. Denote the single relation symbol in \mathcal{R} by R . It holds that Δ is equivalent to a pair of keys and therefore we can denote it by $\{X \rightarrow \llbracket R \rrbracket, Y \rightarrow \llbracket R \rrbracket\}$ where X, Y are subsets of $\llbracket R \rrbracket$. Let $f = R^0(a, b)$. We define a fact-wise reduction $\Pi : R^0 \rightarrow R$, using the constant $\odot \in \text{Const}$. That is, denote $\Pi(f) = (d_1, \dots, d_n)$ where for all $i = 1, \dots, n$

$$d_i \stackrel{\text{def}}{=} \begin{cases} a & i \in X \setminus Y \\ b & i \in Y \setminus X \\ \odot & \text{otherwise.} \end{cases}$$

In order to prove that Π is indeed a fact-wise reduction, we should show it is well-defined, preserves consistency and inconsistency and is injective.

It is straightforward to see that Π is injective. Note that since Δ is not equivalent to a single FD, it holds that $X \setminus Y \neq \emptyset$ and $Y \setminus X \neq \emptyset$. Therefore, there exist i and j such that $d_i = a$ and $d_j = b$. To show that it preserves consistency, we should show that for every two facts f and f' , the set $\{f, f'\}$ is consistent w.r.t Δ^0 if and only if $\{\Pi(f), \Pi(f')\}$ is consistent w.r.t Δ .

The “if” direction. Assume that $\{f, f'\}$ is consistent w.r.t Δ^0 , we contend that $\{\Pi(f), \Pi(f')\}$ is consistent w.r.t Δ . Since $\{f, f'\}$ is consistent w.r.t Δ^0 , if f and f' agree on A then they must also agree on B and if they do not agree on A then they also do not agree on B . Thus, either $a = a', b = b'$ or $a \neq a', b \neq b'$.

If $a = a'$ and $b = b'$ then by Π 's definition, it holds that $\Pi(f) = \Pi(f')$ and thus $\{\Pi(f), \Pi(f')\} = \{\Pi(f)\}$ is consistent w.r.t Δ .

Let $\Pi(f) = (d_1, \dots, d_n)$ and let $\Pi(f') = (d'_1, \dots, d'_n)$. If $a \neq a'$ and $b \neq b'$ then by Π 's definition and since $X \setminus Y$ and $Y \setminus X$ are not empty it holds that there exists $i \in X \setminus Y$ such that $d_i \neq d'_i$ and $j \in Y \setminus X$ such that $d_j \neq d'_j$. Thus $\Pi(f)$ and $\Pi(f')$ do not agree on X nor on Y . That is, $\{\Pi(f), \Pi(f')\}$ is consistent with respect to Δ .

The “only if” direction. Assume $\{f, f'\}$ is inconsistent w.r.t Δ^0 , we contend that $\{\Pi(f), \Pi(f')\}$ is inconsistent w.r.t Δ . Since $\{f, f'\}$ is inconsistent w.r.t Δ^0 , if f and f' agree on A then they must disagree on B and if they disagree on A then they must agree on B . Thus, either $a = a', b \neq b'$ or $a \neq a', b = b'$. Both cases are symmetric and thus we will prove the claim only for the case where $a = a'$ and $b \neq b'$. Let $\Pi(f) = (d_1, \dots, d_n)$ and $\Pi(f') = (d'_1, \dots, d'_n)$. By Π 's definition, $d_i = d'_i$ for $i \in X \setminus Y$ and $d_j \neq d'_j$ for $j \in Y \setminus X$. Moreover, for all other k , it holds that $d_k = d'_k = \odot$. That is, the facts $\Pi(f)$ and $\Pi(f')$ agree on X and disagree on Y . Hence, $\{\Pi(f), \Pi(f')\}$ is inconsistent w.r.t Δ \square

A.4 Proof of Lemma 5.10

LEMMA 5.10. *For all $i = 1, \dots, 5$ there is a fact-wise reduction from \mathbf{S}^0 to \mathbf{S}^i .*

PROOF. Recall our schemas \mathbf{S}^i , for $i = 1, \dots, 5$, where each \mathbf{S}^i is the schema $(\mathcal{R}^i, \Delta^i)$, and \mathcal{R}^i is the singleton $\{R^i\}$. The specification of the \mathbf{S}^i is as follows.

0. $R^0/2$ and $\Delta^0 = \{A \rightarrow B, B \rightarrow A\}$
1. $R^1/3$ and $\Delta^1 = \{AB \rightarrow C, BC \rightarrow A, AC \rightarrow B\}$
2. $R^2/3$ and $\Delta^2 = \{A \rightarrow B, B \rightarrow A\}$
3. $R^3/3$ and $\Delta^3 = \{AB \rightarrow C, C \rightarrow B\}$
4. $R^4/3$ and $\Delta^4 = \{A \rightarrow B, B \rightarrow C\}$
5. $R^5/3$ and $\Delta^5 = \{A \rightarrow C, B \rightarrow C\}$

Let $f = R^0(a, b)$ be a fact over $R^0/2$. For all $i = 1, \dots, 5$ we define the fact-wise reduction $\Pi^i : R^0 \rightarrow R^i$, using $\otimes \in \text{Const}$, by:

1. $\Pi^1(f) \stackrel{\text{def}}{=} R^1(a, b, \otimes)$,
2. $\Pi^2(f) \stackrel{\text{def}}{=} R^2(a, b, \otimes)$,
3. $\Pi^3(f) \stackrel{\text{def}}{=} R^3(\otimes, a, b)$,
4. $\Pi^4(f) \stackrel{\text{def}}{=} R^4(a, b, a)$,
5. $\Pi^5(f) \stackrel{\text{def}}{=} R^5(a, b, \langle a, b \rangle)$.

Regarding the fact-wise reduction Π^5 , note that $\langle a, b \rangle = \langle a', b' \rangle$ if and only if $a = a'$ and $b = b'$. Note that for each $i = 1, \dots, 5$, we have that Π^i is computable in polynomial time. Moreover, it is straightforward to see that for each $i = 1, \dots, 5$, it holds that Π^i is injective. Thus, it is left to show that for each $i = 1, \dots, 5$, the fact-wise reduction Π^i preserves consistency and inconsistency. That is, given two facts $f = R^0(a, b)$ and $f' = R^0(a', b')$, we show that for each $i = 1, \dots, 5$, it holds that $\{\Pi^i(f), \Pi^i(f')\}$ is consistent w.r.t Δ^i , if and only if $\{f, f'\}$ is consistent w.r.t Δ^0 .

We prove the above for $i = 5$.

The “if” direction. Let $f = R^0(a, b)$ and $f' = R^0(a', b')$. Assume that $\{f, f'\}$ is inconsistent with respect to Δ^0 . We contend that $\{\Pi^5(f), \Pi^5(f')\}$ is inconsistent w.r.t Δ^1 . Note that since $\{f, f'\}$ is inconsistent with respect to Δ^0 , f and f' must agree on one attribute and disagree on the second. Both cases are symmetric and thus it suffices to show that if $f = R^0(a, b)$ and $f' = R^0(a', b')$ agree on A and disagree on B then $\{\Pi^5(f), \Pi^5(f')\}$ is inconsistent w.r.t Δ^1 . Indeed, $f = R^0(a, b)$ and $f' = R^0(a', b')$ agree on A which implies $a = a'$, and disagree on B which implies $b \neq b'$. Thus, $\Pi^5(f) = R^1(a, b, \langle a, b \rangle)$ and $\Pi^5(f') = R^1(a', b', \langle a', b' \rangle)$ agree on A but do not agree on C . That is, $\{\Pi^5(f), \Pi^5(f')\}$ is inconsistent w.r.t Δ^1 .

The “only if” direction. Let $f = R^0(a, b)$ and $f' = R^0(a', b')$ and assume that $\{f, f'\}$ is consistent with respect to Δ^0 . We contend that $\{\Pi^5(f), \Pi^5(f')\}$ is consistent w.r.t Δ^1 . If f and f' agree on A (i.e., $a = a'$), since they are consistent w.r.t Δ^0 they must agree also on B (i.e., $b = b'$). Hence $f = f'$ and therefore $\Pi^5(f) = \Pi^5(f')$. We conclude that $\{\Pi^5(f), \Pi^5(f')\}$ is consistent w.r.t Δ^1 . If f and f' do not agree on A (i.e., $a \neq a'$), since they are consistent w.r.t Δ^0 they must also disagree on B (i.e., $b \neq b'$). Thus, $\Pi^5(f)$ and $\Pi^5(f')$ do not agree on A nor on B . That is, Δ^1 holds and thus $\{\Pi^5(f), \Pi^5(f')\}$ is consistent w.r.t Δ^1 .

Similarly, this can be shown for each $i = 1, \dots, 4$ (which are simpler cases than $i = 5$). \square

B. PROOFS FOR SECTION 6

In the current section we prove the correctness of the CCategorycity algorithm, introduced in Section 6. In particular, we prove Theorem 6.5. For convenience, we repeat the theorem here.

THEOREM 6.5. *Let (I, \mathcal{H}, \succ) be an inconsistent prioritizing instance, and let J be the subinstance of I constructed in the execution of CCategorycity(I, \mathcal{H}, \succ). Then the following are equivalent.*

1. J is consistent.
2. There is a single c -repair.

Moreover, if J is consistent then J is the single c -repair.

We will divide our proof into two parts. First we will prove that the algorithm is sound (i.e., if J is consistent, then I has precisely one c -repair). Later, we will prove that the algorithm is complete (i.e., if I has precisely one c -repair, then J is consistent). Before that, we need a basic lemma that will be used in both parts of the proof.

B.1 Basic Lemma

Let (I, \mathcal{H}, \succ) be an inconsistent prioritizing instance over a signature \mathcal{R} . We start by proving the following lemma.

LEMMA B.1. *Suppose that P_1, \dots, P_t are the positive strata constructed by executing $\text{CCategoricity}(I, \mathcal{H}, \succ)$. For every $k \in \{1, \dots, t\}$, if $J_k = P_1 \cup \dots \cup P_k$ is consistent, then there exists an execution of the FindCRep algorithm on (I, \mathcal{H}, \succ) , such that at the beginning of some iteration of that execution the following hold:*

- The set of facts included in J is $P_1 \cup \dots \cup P_k$.
- Every fact $f \in P_{k+1}$ belongs to $\max_{\succ}(I)$.

PROOF. Let us start the execution of the FindCRep algorithm on (I, \mathcal{H}, \succ) as follows. First, we will select all of the facts from P_1 and add them to J (one by one), then we will select all of the facts from P_2 and add them to J and so on. If we can add all of the facts from $P_1 \cup \dots \cup P_k$ to J during this process, and at the end all of the facts in P_{k+1} belong to $\max_{\succ}(I)$, then there exists an execution of the FindCRep algorithm that satisfies all the conditions and that will conclude our proof.

Otherwise, one of the following holds.

- One of the facts in $P_1 \cup \dots \cup P_k$ cannot be added to J .
- After adding all of the facts from $P_1 \cup \dots \cup P_k$ (and only these facts) to J , one of the facts in P_{k+1} does not belong to $\max_{\succ}(I)$.

If the first case holds, let P_i be the first positive stratum such that a fact $f \in P_i$ cannot be added to J after adding all of the facts in $P_1 \cup \dots \cup P_{i-1}$ and maybe some of the facts in P_i to J . Note that $P_1 \cup \dots \cup P_k$ does not contain a hyperedge, thus if f cannot be added to J , it holds that f does not belong to $\max_{\succ}(I)$. That is, there exists at least one fact g , such that $g \succ f$. Moreover, there exists at least one finite sequence, $h_1 \succ \dots \succ h_m \succ f$, of facts in I (with $m \geq 1$), such that $h_1 \in \max_{\succ}(I)$. Note that for each such sequence of facts, it holds that $h \succ^+ f$ for every fact h in the sequence. In this case, one of the following holds for each fact h :

- h belongs to P_j or N_j for some $j < i$.
- h belongs to P_j or N_j for some $j \geq i$.

We assumed that we were able to add all the facts from $P_1 \cup \dots \cup P_{i-1}$ to J , thus $h \in \{h_1, \dots, h_m\}$ cannot belong to P_j for some $j < i$. Furthermore, by the definition of P_i , after removing all of the facts from $P_1 \cup \dots \cup P_{i-1}$ and $N_1 \cup \dots \cup N_{i-1}$, every fact $f \in P_i$ belongs to $\max_{\succ}^+(I)$ (and consequently to $\max_{\succ}(I)$). This cannot be the case if there exists a fact $h \succ^+ f$ that belongs to P_j or N_j for some $j \geq i$. Thus, the only possibility left is that every fact $h \in \{h_1, \dots, h_m\}$ belongs to N_j for some $j < i$.

Since for every such sequence of facts, h_1 belongs to $\max_{\succ}(I)$, it can be selected in line 3 of the algorithm at the next iteration. By the definition of N_j , there exists a hyperedge that is contained in $P_1 \cup \dots \cup P_j \cup \{h_1\}$. We know that all the facts in $P_1 \cup \dots \cup P_j$ were added to J , thus h_1 will be excluded from J . After selecting fact h_1 from each sequence and removing it from I , the next fact, h_2 , in each sequence belongs to $\max_{\succ}(I)$. The previous arguments hold for h_2 as well, thus h_2 can be selected in line 3 of the algorithm at the next iteration. Note that if a fact h belongs to more than one sequence, it will be removed after all the previous facts in each one of these sequences are removed. We can continue with this process until it holds that $f \in \max_{\succ}(I)$, and then add it to J , in contradiction to our assumption.

If the second case holds, at each of the next iterations of the algorithm (and before adding facts that do not belong to $P_1 \cup \dots \cup P_k$ to J), there exists a fact $f \in P_{k+1}$ that does not belong to $\max_{\succ}(I)$. Similarly to the previous part, if f does not belong to $\max_{\succ}(I)$, there exist at least one sequence $h_1 \succ \dots \succ h_m \succ f$ of facts, such that $h \succ^+ f$ for every fact in the sequence and each fact h belongs to N_j for some $j < k+1$. We again can select all of the facts in each sequence as the maximal elements in line 3 of the algorithm by topological order until $f \in \max_{\succ}(I)$. This is a contradiction to our assumption, thus f does belong to $\max_{\succ}(I)$ after adding all of the facts in $P_1 \cup \dots \cup P_k$ to J and removing all of the corresponding facts in $N_1 \cup \dots \cup N_k$ from I . Note that in this case, we will be able to select f in line 3 of the algorithm at the next iteration; however, we may not be able to add f to J , since J_{k+1} is not consistent. \square

EXAMPLE B.2. Consider the CCategoricity execution on the inconsistent prioritizing instance (I, \mathcal{H}, \succ) from our followers running example, illustrated in Figure 6. In this example, $P_1 \cup P_2$ is consistent. We can start building the repair J , using the

FindCRep algorithm as follows. (We recall that we denote inclusion in J by plus and exclusion from J by minus.)

$$+f_{11}, +f_{34}, -f_{12}, -f_{21}, -f_{31}, +f_{22}, +f_{35}$$

As Lemma B.1 states, at the current iteration of the algorithm, all the facts in $P_1 \cup P_2$ belong to J . Moreover, the facts in P_3 (that is, f_{23} and f_{32}) belong to $\max_{\succ}(I)$, as the second part of the lemma states.

Note that in our example, there exists a different execution of the FindCRep algorithm.

$$+f_{11}, -f_{12}, -f_{21}, -f_{31}, +f_{22}, \\ +f_{23}, +f_{32}, -f_{24}, +f_{34}, +f_{35}$$

In this case, there does not exist an iteration of the algorithm that satisfies the conditions from Lemma B.1. Thus, different executions of the algorithm may be possible; however, there always exists an execution of the FindCRep algorithm that satisfies the conditions from the lemma. \square

B.2 Soundness

Next, we prove that if J is consistent, then I has precisely one c-repair. Let us denote by t the number of iterations of the CCategorycity algorithm on the input (I, \mathcal{H}, \succ) . That is, it holds that $J = P_1 \cup \dots \cup P_t$ at the end of the algorithm. Since J is consistent, Lemma B.1 and the fact that the FindCRep algorithm is sound (Theorem 4.1) imply that there exists a c-repair K that includes all of the facts in $P_1 \cup \dots \cup P_t$. The following lemma proves that a c-repair cannot include any fact that belongs to N_i for some $i \in \{1, \dots, t\}$.

LEMMA B.3. *If J is consistent, then no c-repair contains any fact from $N_1 \cup \dots \cup N_t$.*

PROOF. Let us assume, by way of contradiction, that there exists a c-repair K' that includes at least one fact from $N_1 \cup \dots \cup N_t$. By Theorem 4.1, FindCRep is complete, thus there exists an execution of the algorithm that produces K' . Consider an execution of FindCRep that produces K' . In that execution, consider the first time that a fact from a negative stratum is added to K' ; let g be that fact. That is, $g \in N_i$ for some $i \in \{1, \dots, t\}$. By the definition of N_i , there exists a hyperedge $e = \{f_1, \dots, f_m, g\}$ that is contained in $P_1 \cup \dots \cup P_i \cup \{g\}$, such that for every other fact $f \in e$ it holds that $f \succ^+ g$. If we are able to choose fact g for the repair at some iteration of the algorithm, it necessarily belongs to $\max_{\succ}(I)$ at this iteration. Therefore, all of the other facts in e are no longer included in I .

This may be the case if all of the other facts in the hyperedge were already added to K' . However, adding g to K' will result in a hyperedge, in contradiction to the fact the K' is a repair. Therefore, at least one of the facts in the hyperedge was removed from I without being added to K' ; let f be such a fact. That is, there exists another hyperedge e' that includes f , such that all of the other facts in this hyperedge were added to K' before f was removed from I . All of these facts, including f , belong to $P_1 \cup \dots \cup P_t$, since we assumed that g is the first fact from some N_i that was chosen for the repair. Hence, we found a hyperedge, e' , that is contained in $P_1 \cup \dots \cup P_t$, in contradiction to the fact that $P_1 \cup \dots \cup P_t$ is consistent. Thus, a fact $g \in N_i$ cannot be added to any repair of I . \square

Lemma B.3 implies that every c-repair is contained in J . Moreover, as said above, there exists a c-repair that includes all of the facts in J . Thus, J is the only c-repair of I and this concludes the proof of soundness of the CCategorycity algorithm.

B.3 Completeness

Finally, we prove that if I has precisely one c-repair, then J is consistent. Throughout this section we fix (I, \mathcal{H}, \succ) and assume that there is exactly one c-repair, which we denote by K . Let J_i denote the subinstance $P_1 \cup \dots \cup P_i$. We prove by induction on k that after the k th iteration of CCategorycity, the instance $J_k = P_1 \cup \dots \cup P_k$ is consistent. Then, we will conclude that J is consistent when the algorithm reaches line 9, and consequently, the algorithm returns true as we expect.

The basis of the induction, $k = 0$, is proved by observing that J_0 is an empty set, thus it does not include any hyperedge. For the inductive step, we need to prove that if J_k is consistent then J_{k+1} is also consistent. So, suppose that J_k is consistent. Let us assume, by way of contradiction, that J_{k+1} is inconsistent, that is, J_{k+1} contains a hyperedge. We next prove the following lemma.

LEMMA B.4. *There exists a fact $f \in P_{k+1}$ and a hyperedge e such that the following hold:*

- $f \in e$.
- $(e \setminus \{f\}) \subseteq P_1 \cup \dots \cup P_k$.

PROOF. Let e be a hyperedge that is contained in $P_1 \cup \dots \cup P_{k+1}$ and has a minimal intersection with P_{k+1} . Let $\{f_1, \dots, f_m\}$ be the set $e \cap P_{k+1}$. Then $m > 0$ since J_k does not contain a hyperedge. To prove the lemma, we need to show that $m = 1$. Suppose, by way of contradiction, that $m > 1$. Since J_k is consistent, Lemma B.1 implies that we can start building a c-repair using the FindCRep algorithm by first choosing all of the facts in $P_1 \cup \dots \cup P_k$. Moreover, after choosing all of these facts, there exists an iteration i in which each fact in P_{k+1} belongs to $\max_{\succ}(I)$. Since FindCRep always produces a c-repair

(Theorem 4.1), we can now choose the fact f_1 , which will result in a c-repair J_1 . This holds true due to our assumption that $m > 1$ and m is minimal (hence, adding a single fact to a set of facts that currently includes only facts from $P_1 \cup \dots \cup P_k$ does not result in the containment of a hyperedge).

The c-repair J_1 cannot contain all of the facts in $\{f_1, \dots, f_m\}$, since a repair cannot contain a hyperedge. Let us assume that $f_j \in \{f_1, \dots, f_m\}$ is not in J_1 . If we choose f_j instead of f_1 at the i th iteration, the FindCRep algorithm will produce a c-repair J_2 that includes f_j . Again, we can choose f_j due to our assumption that $m > 1$ and m is minimal. That is, we have two distinct c-repairs, J_1 and J_2 , in contradiction to our assumption that I has precisely one c-repair. \square

Since we assumed that I has exactly one c-repair and $P_1 \cup \dots \cup P_k$ is consistent, Lemma B.4 implies that there exists a hyperedge $e = \{f_1, \dots, f_m\}$ such that precisely one of the f_i belongs to P_{k+1} , while the other facts belong to $P_1 \cup \dots \cup P_k$. Without loss of generality, we can assume that $f_1 \in P_{k+1}$. We next prove the existence of two distinct c-repairs:

- A c-repair that includes all of the facts in $e \setminus \{f_1\}$ and does not include f_1 .
- A c-repair that includes f_1 .

LEMMA B.5. *There exists a c-repair that does not include f_1 .*

PROOF. Lemma B.1 implies that it is possible to build a c-repair using FindCRep by first choosing all of the facts in $P_1 \cup \dots \cup P_k$. As all of the facts in $\{f_2, \dots, f_m\}$ belong to $P_1 \cup \dots \cup P_k$, all of them will be chosen in this process as well. Since the FindCRep algorithm is sound (Theorem 4.1), this specific execution of the algorithm will result in a c-repair J of I that includes all of the facts f_2, \dots, f_m . The fact f_1 cannot be included in J , since adding it to J will result in a repair that contains a hyperedge, which is impossible by definition. \square

To complete the proof of completeness, we prove that there exists another c-repair of I that includes f_1 . In order to do so, we again take advantage of the algorithm FindCRep, and prove the following lemma.

LEMMA B.6. *There exists a c-repair that includes f_1 .*

PROOF. In order to prove the lemma, we start building the corresponding c-repair, using the FindCRep algorithm, by first selecting in line 3 of the algorithm only facts from $P_1 \cup \dots \cup P_k$ that either cannot be added to J or satisfy at least one of the following conditions:

- $g \succ^+ f_1$.
- f_1 and g are not neighbors in \mathcal{H} .

Note that not all of the facts in $P_1 \cup \dots \cup P_k$ that satisfy at least one of the conditions can be selected in this process. A fact $g \in P_2$, for example, may be left out of J if a fact $f \in P_1$ is not selected because it does not satisfy any of the conditions and it holds that $f \succ g$. In this case, g will not belong to $\max_{\succ}(I)$ until f is selected by the algorithm. However, since J_k is consistent, all of the facts from $P_1 \cup \dots \cup P_k$ that can be selected in line 3 of the algorithm during this process can also be added to J .

We will now prove that after selecting all of these facts, the algorithm can add the fact f_1 to J next (that is, before adding any other fact to J). This will eventually result in a c-repair that includes f_1 and will conclude our proof. Let us assume, by way of contradiction, that we cannot add f_1 to J next. That is, one of the following holds:

- There exists a hyperedge that contains f_1 such that all of the other facts in the hyperedge have already been added to J .
- We must add another fact to J before it holds that $f_1 \in \max_{\succ}(I)$.

If the first case holds, then there exists a hyperedge e , such that $f_1 \in e$ and all of the other facts in the hyperedge were added to J in the previous iterations of the algorithm. Note that all of the facts that have already been added to J , including $e \setminus \{f_1\}$, belong to $P_1 \cup \dots \cup P_k$. Moreover, it holds that $h \succ^+ f_1$ for every fact $h \in e \setminus \{f_1\}$, since a fact $h \in e \setminus \{f_1\}$ that does not hold $h \succ^+ f_1$, does not satisfy any of the conditions and could not have been added to J . Hence, we found a hyperedge e , such that $f_1 \in e$, and for every other fact $h \in e$ it holds that $h \in P_1 \cup \dots \cup P_k$ and $h \succ^+ f_1$. By the definition of negative stratum, it should hold that $f_1 \in N_i$ for some $i \in \{1, \dots, k+1\}$ (the exact value of i depends on the other facts in this hyperedge), in contradiction to the fact that $f_1 \in P_{k+1}$. Thus, the first case is impossible.

If the second case holds, then at one of the next iterations, only facts that can be added to J belong to $\max_{\succ}(I)$, while f_1 does not belong to $\max_{\succ}(I)$ yet. In this case, the FindCRep algorithm must add another fact to J before adding f_1 to J . Since f_1 does not belong to $\max_{\succ}(I)$, there exists at least one fact g such that $g \succ f_1$. Moreover, there exists at least one finite sequence, $h_1 \succ \dots \succ h_m \succ f$, of facts in I (with $m \geq 1$), such that $h_1 \in \max_{\succ}(I)$. Note that for each such sequence of facts, it holds that $h \succ^+ f$ for every fact h in the sequence. In this case, one of the following holds for each fact h :

- h belongs to P_j or N_j for some $j < k+1$.
- h belongs to P_j or N_j for some $j \geq k+1$.

By the definition of P_{k+1} , after removing all of the facts from $P_1 \cup \dots \cup P_k$ and $N_1 \cup \dots \cup N_k$, every fact $f \in P_{k+1}$ belongs to $\max_{\succ}^+(I)$. This cannot be the case if there exists a fact $h \succ^+ f$ that belongs to P_j or N_j for some $j \geq k+1$. Thus, each fact h either belongs to P_j or N_j for some $j < k+1$. If h_1 belongs to some P_j , then h_1 was not added to J since it does not satisfy any of the conditions. This holds true since $h_1 \in \max_{\succ}^+(I)$, thus it can be selected in line 3 of the algorithm. That is, there exists a hyperedge $e \subseteq P_1 \cup \dots \cup P_{k+1}$, such that $\{f_1, h_1\} \subseteq e$ and it does not hold that $h_1 \succ^+ f_1$. This is a contradiction to the fact that $h_1 \succ^+ f$, thus this cannot be the case.

The only possibility left is that $h_1 \in N_j$ for some $j < k+1$. Since it also holds that $h_1 \in \max_{\succ}^+(I)$, the fact h_1 can be selected in line 3 of the algorithm next. Note that by the definition of N_j , there exists a hyperedge that is contained in $P_1 \cup \dots \cup P_j \cup \{h_1\}$, such that for every other fact h' in the hyperedge it holds that $h' \succ^+ h_1$. Since $h_1 \in \max_{\succ}^+(I)$ none of these facts still belongs to I , thus they have already been added to J , and h_1 will not be added to J , but only removed from I . We can continue with this process and choose all of the facts in each one of the corresponding sequences by topological order, until $f_1 \in \max_{\succ}^+(I)$. (If a fact h belongs to more than one sequence, it will be removed after all the previous facts in each one of these sequences are removed.)

Then, we can add f_1 to J , since every hyperedge that contains f_1 and is included in $P_1 \cup \dots \cup P_k \cup \{f_1\}$ also contains at least one fact g that does not satisfy $g \succ^+ f$ (otherwise, f_1 would belong to some N_j). That is, each hyperedge that contains f_1 contains at least one fact that does not satisfy any of the conditions and was not added to J , and adding f_1 to J will not close any hyperedge. This is a contradiction to our assumption, thus f_1 can be chosen for the repair next. As said above, this concludes our proof. \square

From Lemma B.5 and Lemma B.6 we can conclude that there exist two distinct c-repairs of I . This is a contradiction to our assumption that I has precisely one c-repair. Hence, J_{k+1} is necessarily consistent, and this concludes our proof of completeness.

C. PROOFS FOR SECTION 7

C.1 Proof of Theorem 7.2

THEOREM 7.2. $g\text{-categoricity}(\mathbf{S}^6)$ is Π_2^P -complete.

PROOF. To show Π_2^P -hardness, we construct a reduction from QCNF_2 to the problem of $g\text{-categoricity}(\mathbf{S}^6)$. The input to QCNF_2 consists of a CNF formula $\psi(\mathbf{x}, \mathbf{y})$ where \mathbf{x} and \mathbf{y} are disjoint sequences of variables. The goal is to determine whether for every assignment to \mathbf{x} there exists an assignment to \mathbf{y} such that the two satisfy ψ . We denote \mathbf{x} by x_1, \dots, x_n , \mathbf{y} by y_1, \dots, y_k and $\psi = c_1 \wedge \dots \wedge c_m$. The input (I, \succ) is constructed from ψ by adding to it the following facts:

- $R^6(0, x_i, 0)$ and $R^6(0, x_i, 1)$ for each variable x_i ,
- $R^6(1, x_i, 0)$ and $R^6(1, x_i, 1)$ for each variable x_i ,
- $R^6(1, y_l, 0)$ and $R^6(1, y_l, 1)$ for each variable y_l ,
- $R^6(0, c_j, c_j)$ for each clause c_j
- $R^6(2, 0, 0)$ which we denote by f_0 .

The priority is defined by:

- $R^6(1, x_i, b) \succ R^6(0, x_i, b)$ for all $b = 0, 1$ and for all variables x_i ,
- $R^6(1, x_i, 1) \succ R^6(0, c_j, c_j)$ if x_i appears in clause c_j ,
- $R^6(1, x_i, 0) \succ R^6(0, c_j, c_j)$ if $\neg x_i$ appears in clause c_j and
- $f_0 \succ R^6(1, w, b)$ for all variables w and $b = 0, 1$.

Our construction is illustrated in Figure 11 for the following input to QCNF_2 : $\mathbf{x} = x, w$, $\mathbf{y} = y, z$ and $\psi = (w \vee y \vee z) \wedge (x \vee \neg w \vee y) \wedge (\neg x \vee \neg w \vee \neg z)$.

Observe that the subinstance $\{f_0\}$ is a c-repair and thus it is also a g-repair. To complete the proof we will show that $\{f_0\}$ is the only g-repair if and only if ψ is a “yes” instance; that is, for every assignment to \mathbf{x} there exists an assignment to \mathbf{y} such that the two satisfy ψ .

The “if” direction.: Assume that ψ is a “yes” instance and let J be a g-repair of I . We contend that J is exactly $\{f_0\}$. Assume, by way of contradiction that $J \neq \{f_0\}$. Since J is consistent and since Δ^6 contains the FD $\emptyset \rightarrow A$, it holds that all of the facts in J agree on A . By our assumption that $J \neq \{f_0\}$, there are two cases:

1. For each fact $f \in J$, it holds that $f[A] = 1$. By \succ ’s definition, we have that $\{f_0\}$ is a global improvement of J which is in contradiction to J being a g-repair.




Figure 11: Illustration of the reduction from QCNF₂ to g-categoricity(S⁶)

2. For each fact $f \in J$, it holds that $f[A] = 0$. We state that for all j , the fact $R^6(0, c_j, c_j) \in J$. Assume, by way of contradiction, there exists j for which $R^6(0, c_j, c_j) \notin J$. Since J is maximal it contains a fact f that is inconsistent with $R^6(0, c_j, c_j)$. Since for each fact $f \in J$, it holds that $f[A] = 0$, there is no such fact and as a consequence no such j . Hence for all j , we have that $R^6(0, c_j, c_j) \in J$. Moreover, since J is maximal, it must contain a fact $R^6(0, x_i, b)$ for each $x_i \in \mathbf{x}$. Note that Δ^6 contains the FD $B \rightarrow C$. This insures that J cannot contain both of the facts $R^6(0, x_i, 0)$ and $R^6(0, x_i, 1)$. This implies that J encodes an assignment $\tau_{\mathbf{x}}$ for the variables $x_i \in \mathbf{x}$. This assignment is given by

$$\tau_{\mathbf{x}}(x_i) \stackrel{\text{def}}{=} \begin{cases} 0 & R^6(0, x_i, 0) \in J \\ 1 & R^6(0, x_i, 1) \in J \end{cases}$$

Since ψ is a “yes” instance, there exists an assignment $\tau_{\mathbf{y}}$ for the variables $y_l \in \mathbf{y}$ that together with $\tau_{\mathbf{x}}$ satisfies ψ . Let K be the subinstance of I that consists of the facts $R^6(1, x_i, \tau_{\mathbf{x}}(x_i))$ for all i and $R^6(1, y_l, \tau_{\mathbf{y}}(y_l))$ for all l . Note that K is consistent. We contend that K is a global improvement of J . Since J and K are disjoint, it suffices to show that for every fact $f \in J$ there is a fact $f' \in K$ such that $f' \succ f$. Let $f \in J$. If f is of the form $R^6(0, x_i, b)$ then we choose $f' = R^6(1, x_i, b)$. If f is of the form $(0, c_j, c_j)$, then we choose $f' = R^6(1, w, b)$ where w is the variable of a literal that satisfies c_j under the union of $\tau_{\mathbf{x}}$ and $\tau_{\mathbf{y}}$. We conclude that K is a global improvement of J . That is, J is not a g-repair in contradiction to our assumption.

The “only if” direction.: Assume that $\{f_0\}$ is the only g-repair of I . We contend that for every assignment to \mathbf{x} there exists an assignment to \mathbf{y} such that the two satisfy ψ . Let $\tau_{\mathbf{x}}$ be an assignment for \mathbf{x} and let J be a consistent subinstance of I that consists of the facts $R^6(0, c_j, c_j)$ for $j = 1, \dots, k$ and $R^6(0, x_i, \tau_{\mathbf{x}}(x_i))$ where $x_i \in \mathbf{x}$. Since $\{f_0\}$ is the only g-repair, it holds that J has a global improvement. Let us denote such a global improvement by K . Assume, without loss of generality, that K is maximal (if not, we can extend K with additional facts). By \succ ’s definition, K must consist of facts of the form $R^6(1, w, b)$. Since the FD $\emptyset \rightarrow A$ is in Δ^6 , for each fact $f \in K$, it holds that $f[A] = 1$. Since K is maximal and since $B \rightarrow C$ is in Δ^6 , we have that K encodes an assignment τ for \mathbf{x} and \mathbf{y} . This assignment is given by

$$\tau(w) \stackrel{\text{def}}{=} \begin{cases} 0 & R^6(1, w, 0) \in K \\ 1 & R^6(1, w, 1) \in K \end{cases}$$

Since K is a global improvement of J , it must contain the fact $R^6(1, x_i, b_i)$ whenever $R^6(0, x_i, b_i)$ is in J (this is true since no other fact has a priority over $R^6(0, x_i, b_i)$). Therefore τ extends $\tau_{\mathbf{x}}$. Finally we observe that every c_j is satisfied by τ . A satisfying literal is one that corresponds to a fact $R^6(1, w, b)$ that satisfies $R^6(1, w, b) \succ c_j$. We conclude that ψ is a “yes” instance as claimed. \square

C.2 Proof of Theorem 7.3

THEOREM 7.3. *Let $\mathbf{S} = (\mathcal{R}, \Delta)$ be a schema such that \mathcal{R} consists of a single relation symbol R and Δ consists of two nontrivial FDs $X \rightarrow Y$ and $W \rightarrow Z$. Suppose that each of W and Z contains an attribute that is in none of the other three sets. Then g-categoricity(\mathbf{S}) is Π_2^P -complete.*

PROOF. Recall the schema $\mathbf{S}^6 = (\mathcal{R}, \Delta)$ where \mathcal{R} consists of a single ternary relation $R^6/3$ and $\Delta^6 = \{\emptyset \rightarrow A, B \rightarrow C\}$. We define a fact-wise reduction $\Pi : R^6 \rightarrow R$, using the constants $\odot, \oplus \in \text{Const}$. Let $f = R^6(a, b, c)$. We define Π by $\Pi(f) = R^6(d_1, \dots, d_n)$ where for all $i = 1, \dots, n$

$$d_i \stackrel{\text{def}}{=} \begin{cases} \odot & i \in X \\ a & i \in Y \setminus X \\ b & i \in W \setminus (X \cup Y \cup Z) \\ c & i \in Z \setminus (X \cup Y \cup W) \\ \oplus & \text{otherwise} \end{cases}$$

It is left to show that Π is a fact-wise reduction. To do so, we prove that Π is well defined, is injective and preserves consistency and inconsistency.

Π is well defined. It suffices to show that each d_i is well-defined. We show that the sets in the definition of d_i are pairwise disjoint. Indeed, X is disjoint from the sets $Y \setminus X$, $W \setminus (X \cup Y \cup Z)$ and $Z \setminus (X \cup Y \cup W)$. Moreover, $Y \setminus X$ is a subset of Y and therefore it is disjoint from $W \setminus (X \cup Y \cup Z)$ and $Z \setminus (X \cup Y \cup W)$. Clearly, $W \setminus (X \cup Y \cup Z)$ and $Z \setminus (X \cup Y \cup W)$ are disjoint. Hence, each d_i is well-defined.

Π is injective. Let $f, f' \in R^6$ where $f = R^6(a, b, c)$ and $f' = R^6(a', b', c')$. Assume that $\Pi(f) = \Pi(f')$. Let us denote $\Pi(f) = (d_1, \dots, d_n)$ and $\Pi(f') = (d'_1, \dots, d'_n)$. Note that $Y \setminus X$ is not empty since the FD $X \rightarrow Y$ is nontrivial. Moreover, $W \setminus (X \cup Y \cup Z)$ and $Z \setminus (X \cup Y \cup W)$ are not empty since each of W and Z contains an attribute that is in none of the other three sets. Therefore, there are i, j and k such that $d_i = a$, $d_j = b$ and $d_k = c$. Hence, $\Pi(f) = \Pi(f')$ implies that $d_i = d'_i$, $d_j = d'_j$ and $d_k = d'_k$. Therefore we obtain $a = a'$, $b = b'$ and $c = c'$ which implies $f = f'$.

Π preserves consistency. Let $f = R^6(a, b, c)$ and $f' = R^6(a', b', c')$. We contend that $\{f, f'\}$ is consistent w.r.t Δ^6 if and only if $\{\Pi(f), \Pi(f')\}$ is consistent w.r.t Δ .

The “if” direction:

Assume $\{f, f'\}$ is consistent w.r.t Δ^6 . We prove that $\{\Pi(f), \Pi(f')\}$ is consistent w.r.t Δ . Note that $\Pi(f)$ and $\Pi(f')$ agree on X since for each $i \in X$ we have that $d_i = \odot$, regardless of the input. Since $\{f, f'\}$ is consistent w.r.t Δ^6 , it holds that $a = a'$. By the definition of Π and since $a = a'$, we have that $\Pi(f)$ and $\Pi(f')$ agree on Y . Hence, $\{\Pi(f), \Pi(f')\}$ satisfies the constraint $X \rightarrow Y$. Assume that $\Pi(f)$ and $\Pi(f')$ agree on W . By the definition of Π , since $W \setminus (X \cup Y \cup Z)$ is not empty, it holds that $b = b'$. Since $\{f, f'\}$ is consistent w.r.t Δ^6 , the fact that $b = b'$ implies that also $c = c'$ (due to the constraint $B \rightarrow C$). Hence, $\Pi(f) = \Pi(f')$. This implies and that $\{\Pi(f), \Pi(f')\} = \{\Pi(f)\}$ is consistent w.r.t Δ .

The “only if” direction:

Assume $\{f, f'\}$ is inconsistent w.r.t Δ^6 . We prove that $\{\Pi(f), \Pi(f')\}$ is inconsistent w.r.t Δ . There are two cases

- f and f' do not agree on A . It holds, by Π 's definition, that $\Pi(f)$ and $\Pi(f')$ agree on X . Nevertheless, since f and f' do not agree on A , we have that $a \neq a'$. Hence $\Pi(f)$ and $\Pi(f')$ do not agree on Y . That is, the constraint $X \rightarrow Y$ is not satisfied, which leads us to the conclusion that $\{\Pi(f), \Pi(f')\}$ is inconsistent w.r.t Δ .
- f and f' agree on A . Since $\{f, f'\}$ is inconsistent w.r.t Δ^6 , we have that f and f' agree on B (i.e., $b = b'$) but disagree on C (i.e., $c \neq c'$). Note that $\Pi(f)$ and $\Pi(f')$ agree on W since $a = a'$ and $b = b'$. Nevertheless, they do not agree on Z since $c \neq c'$ and the set $Z \setminus (X \cup Y \cup W)$ is not empty. That is, the constraint $W \rightarrow Z$ is not satisfied which leads us to the conclusion that $\{\Pi(f), \Pi(f')\}$ is inconsistent w.r.t Δ .

□